



Classification of Radar Targets Using Invariant Features

DISSERTATION

Gregory J. Meyer, Captain, USAF

AFIT/DS/ENG/03-04

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this dissertation are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense or the U. S. Government.

CLASSIFICATION OF RADAR TARGETS USING
INVARIANT FEATURES

DISSERTATION

Presented to the Faculty of the Graduate School of Engineering and Management

of the Air Force Institute of Technology

In Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy in Electrical Engineering

Gregory J. Meyer, M.S.E.E.

Captain, United States Air Force

April, 2003

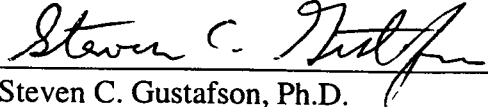
Approved for public release; distribution unlimited

CLASSIFICATION OF RADAR TARGETS USING
INVARIANT FEATURES

Dissertation, April, 2003

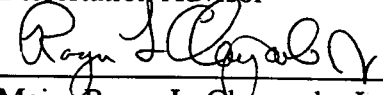
Gregory J. Meyer, M.S.E.E.
Captain, United States Air Force

Approved:



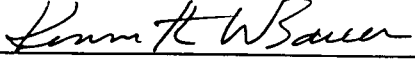
Steven C. Gustafson, Ph.D.
Dissertation Advisor

14 Apr 03
Date



Major Roger L. Claypoole, Jr., Ph.D.
Committee Member

14 APR 03
Date



Kenneth W. Bauer, Ph.D.
Committee Member

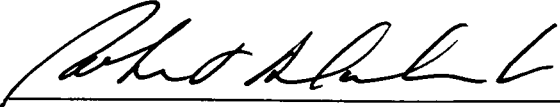
14 APR 03
Date



Mark E. Oxley, Ph.D.
Dean's Representative

14 Apr 03
Date

Accepted:



Robert A. Calico, Jr., Ph.D.
Dean, Graduate School of Engineering and Management

Acknowledgements

I thank my committee members for their time, patience, and understanding during the development of this dissertation. In particular, my research advisor, Dr. Steven Gustafson, acted as mentor and editor for the last three years. I thank my sponsor (AFRL/SNAT) for taking the time to bring me up to speed on invariant theory, especially Dr. Greg Arnold, who kept my research moving in a useful direction. Special thanks go to Dr. Mark Stuff for explaining three-dimensional reconstruction in a way I could understand and who first suggested looking at singular values of object matrices as features for ATR. I thank Dr. Won Roh for deciding that I could handle doctorate level classes and research and for giving me the opportunity to pursue another advanced degree at AFIT.

This dissertation would not be possible without the sponsorship of the Air Force Research Laboratory Sensors Directorate, specifically the Target Recognition Branch (AFRL/SNAT). AFRL/SNAT has long been at the forefront of using invariant techniques to solve SAR ATR problems and is the “center of gravity” of all ATR in the DoD. The groundbreaking effort of Veridian (formerly known as ERIM) in the development of computer algorithms to extract and track three-dimensional scattering centers provided the knowledge and data needed to attack the three-dimensional classification of SAR targets.

Finally, I thank my family for their support and forbearance, especially my wife, who has shown she can raise four children and maintain a household alone. I thank God for everything, with whom anything is possible.

Table of Contents

	Page
Acknowledgements	i
Table of Contents	ii
List of Figures	v
List of Tables	vi
Abstract	vii
Nomenclature	ix
 1 Introduction.....	 1-1
1.1 Background	1-2
1.1.1 Invariance	1-2
1.1.2 Scattering centers.....	1-7
1.1.3 Existing two-dimensional classifiers	1-8
1.2 Three-dimensional MAGI concept	1-9
1.2.1 Three-dimensional reconstruction	1-10
1.2.2 3D MAGI algorithm	1-11
1.3 Problem statement.....	1-14
1.3.1 Three-dimensional classification	1-14
1.3.2 Designing an optimal three-dimensional classifier.....	1-15
 2 Model Building	 2-1
2.1 Three-dimensional scattering model	2-1
2.2 Extraction procedure	2-3
2.2.1 Extracting amplitude and “wrapped” phase.....	2-3
2.2.2 Unwrapping phase	2-4
2.2.3 Solving for scattering center coordinates.....	2-5
2.3 Algorithm validation and test.....	2-7
2.4 Implementation	2-10
2.5 Impact on classifier design.....	2-11

3	Singular Value Features	3-1
3.1	Definition of singular values	3-1
3.2	Invariance of singular values	3-2
3.3	Monomial expansion of object matrix	3-2
3.4	Determining the feature set	3-3
3.5	Feature set definition.....	3-6
4	Classifier Design.....	4-1
4.1	Classifier design and compromises	4-1
4.2	Scattering center stability	4-1
4.3	Target models.....	4-2
4.4	Singular value feature extraction	4-4
4.5	Creating target subclasses	4-6
4.6	Classifier training.....	4-6
4.7	Decision Making	4-8
5	Classifier Test	5-1
5.1	Replicating “real world” conditions.....	5-1
5.2	Calculating confusion matrices	5-2
5.3	Calculating figures of merit	5-4
5.4	Test results	5-5
5.5	Comparisons with existing classifiers.....	5-5
6	Contributions	6-1
6.1	Extract “cloud of points” model.....	6-2
6.2	Generate singular value features	6-2
6.3	Characterize aspect-angle subclasses	6-3
6.4	Find clutter threshold and classify test points	6-3
6.5	Generate confusion matrices	6-4
6.6	Plot PCID and PFID for each target.....	6-4
7	Conclusions and Recommendations	7-1
7.1	Conclusions.....	7-1
7.1.1	Feature selection	7-1
7.1.2	Model building	7-2

7.1.3	Classifier design	7-2
7.1.4	Classifier performance.....	7-2
7.2	Recommendations.....	7-3
7.2.1	Comprehensive testing.....	7-3
7.2.2	Obscured target improvements.....	7-3
7.2.3	Improved classification algorithm with confidence measures.....	7-3
Bibliography.....		BIB-1
Appendix 1.....		A-1
Appendix 2.....		A-12

List of Figures

	Page
Figure 1. Representation of an invariant function.....	1-3
Figure 2. Three-dimensional reconstruction using six.	1-12
Figure 3. Two difficult cases for “unwrapping” phase.	2-6
Figure 4. Unwrapping phase between two scattering centers on a larger object.	2-9
Figure 5. The significance of features from successive monomial expansions.	3-5
Figure 6. Eight military vehicles used as targets.....	4-3
Figure 7. Monomial expansions of the target matrix	4-5
Figure 8. Plot of visible scattering centers as a function of aspect angle.....	4-7
Figure 9. Correlation as a function of aspect angle.....	4-7
Figure 10. Data compression illustration	4-9
Figure 11. Flowchart of the classifier decision process..	4-11
Figure 12. Example of decision and clutter thresholds..	4-12
Figure 13. Plot of PCID and PFID for targets 1 & 2.....	5-6
Figure 14. Plot of PCID and PFID for targets 3 & 4.....	5-7
Figure 15. Plot of PCID and PFID for targets 5 & 6.....	5-8
Figure 16. Plot of PCID and PFID for targets 7 & 8.....	5-9

List of Tables

	Page
Table 1. Relative Motion vs Non-zero Eigenvalues of System	1-11
Table 2. Confusion Matrix for 10% Obscuration	5-3
Table 3. Confusion Matrix for 15% Obscuration	5-3
Table 4. Confusion Matrix for 25% Obscuration	5-3
Table 5. Confusion Matrix for 40% Obscuration	5-3

Abstract

Automatic target recognition (ATR) using radar commonly relies on modeling a target as a collection of point scattering centers. Features extracted from these scattering centers for input to a target classifier may be constructed that are invariant to translation and rotation, i.e., they are independent of the position and aspect angle of the target in the radar scene.

Here an iterative approach for building effective scattering center models is developed, and the shape space of these models is investigated. Experimental results are obtained for three-dimensional scattering centers compressed to nineteen-dimensional feature sets, each consisting of the singular values of the matrix of scattering center locations augmented with the singular values of its second and third order monomial expansions. These feature sets are invariant to translation and rotation and permit the comparison of targets modeled by different numbers of scattering centers. A Mahalanobis distance metric is used that effectively identifies targets under “real world” conditions that include noise and obscuration.

In particular, eight targets of military interest are sampled in tenth-degree aspect angle increments to extract scattering centers, and 36 subclasses that encompass ten degrees are specified for each target. Each subclass is compressed to a nineteen-dimensional singular value feature set, and because the spatial distribution of the 100 nineteen-dimensional points in each subclass is approximately Gaussian, a mean and a covariance matrix represent each subclass. An unknown target is represented as a point in the nineteen-dimensional feature space and matched to the closest subclass mean in Mahalanobis distance (Euclidean distance after the nineteen coordinate axes are rotated

and scaled so that the subclass covariance matrix is the unit matrix). Targets with zero mean Gaussian noise added independently to their scattering centers are matched in both target identity and aspect angle (to within ten degrees) for noise variances up to that of the mean squared deviation of the scattering centers. Also, as demonstrated using plots of probability of correct and false identification, targets obscured by 20% are identified correctly in 80% of the test cases.

This research has developed a practical means (1) to build scattering center models, (2) to compress scattering centers into a small set of invariant features for target classification, and (3) to implement classifiers that effectively function in the presence of noise and obscuration for targets of military interest.

Nomenclature

N	Number of vertices or points in a model
Φ	Invariant function
G	Group of allowable transformations
g	Single element in the group of transformations
I_{xx}	Output of a three-dimensional invariant function
i_x	Output of a one-dimensional invariant function
ρ_x	Coefficient of radar transform from 3D to 1D
z_x^O	Object coordinate
u_x^K	Real part of Kendall coordinate
v_x^K	Imaginary part of Kendall coordinate
d_f	Procrustes distance
μ_x	Mean of a statistical distribution in one dimension
σ_x	Standard deviation of a statistical distribution in one dimension
H_0	Null hypothesis
$F_0(x)$	Postulated distribution function
$F(x)$	Experimental distribution function
A_n	Maximum amplitude of nth scattering center
(x_n, y_n, z_n)	Coordinates of nth scattering center
f	Frequency of swept radar pulse
f_c	Center frequency of swept radar pulse
(ϕ'_n, θ'_n)	Location of nth scattering center maximum on viewing sphere
γ'_n	Exponential damping rate of nth scattering center in phi direction
γ''_n	Exponential damping rate of nth scattering center in theta direction
$\Delta\phi_{\min}$	Minimum sampling interval needed for unambiguous extraction
l_{\max}	Maximum extent of target

Classification of Radar Targets Using Invariant Features

1 Introduction

The Department of Defense has long been interested in performing Automatic Target Recognition (ATR) on radar images, and, more recently, on Synthetic Aperture Radar (SAR) images. Higher resolution sensors can be designed and constructed to improve SAR ATR performance, but for difficult problems (identifying complicated targets of similar type in clutter) this approach provides limited improvement. For example, the effect of resolution on target identification performance using artificially degraded high resolution SAR imagery is quantified in Appendix 1. Using four very similar targets and comparing SAR classification performance over five different resolutions, it was determined that each doubling of resolution generates only a five to ten percent improvement in target identification performance.

Thus, rather than advancing the development of more powerful and expensive sensors, this dissertation examines the use of more powerful techniques to improve the classification performance of existing sensors. Specifically, a classifier that uses invariant features (i.e., inputs which are independent of target translation and rotation) and that uses three-dimensional target information (i.e., information which depends on azimuth, elevation, and range to many scattering center points on the target) is designed and characterized. A successful classifier of this type may change the way future classifiers are developed.

This introductory chapter presents background, including material on invariants, scattering center representations, and existing two-dimensional classifiers; discusses the three-dimensional Motion and Geometric Invariant (3D MAGI) concept; and states the problem addressed by the dissertation. The 3D MAGI algorithm, which was developed elsewhere [48], demonstrates that it is possible to extract target scattering centers which persistent for up to 20 degrees in aspect angle. This persistence and the data compression enabled by the use of the invariant singular value features investigated in this dissertation make a three-dimensional classifier possible.

1.1 Background

This dissertation leverages two technologies: invariance theory and scattering center extraction. Invariance theory provides criteria for identifying useful features from “raw” data and performs exceptionally well with point representations of target objects. The extraction of point-scattering centers from radar images provides useful data compression. The natural marriage of these two technologies enables the effective classification of radar (including SAR) targets.

1.1.1 Invariance

Suppose that an object x in the n -dimensional space \mathcal{R}^n undergoes a group transformation g (g operates on object x denoted by $g(x)$) from a set of allowable group actions G , i.e., $g \in G$ [16, 31]. Letting Φ be any function on \mathcal{R}^n , Figure 1 shows that Φ is an invariant functional if it yields the same value for each object after all possible transformations. That is, $\Phi(g(x)) = \Phi(x)$ for $\forall x \in \mathcal{R}^n$ and $\forall g \in G$. Functions that are nearly

constant after all possible transformations (quasi-invariants) are useful because in many cases they can be treated as invariants [5].

When a three-dimensional object undergoes transformations from the appropriate group action and is projected onto a two-dimensional subspace, useful invariant features can be constructed for the projected object in this subspace. This projection applies to tracking a moving target with SAR: the target undergoes appropriate group actions (translation, rotation, and scale) and is projected to a two-dimensional subspace by the radar.




Figure 1. Representation of an invariant function. Here an object in n-dimensional space undergoes an allowable group action g . A function Φ is invariant if it yields the same value regardless of the group action. In radar ATR allowable group actions typically translate, rotate, or scale the object.

Allowable transformations are equivalent to changing the coordinates (e.g., multiplying by a unitary rotation matrix), so an invariant function is a coordinate-free description of an object: the invariant function is constant with respect to certain coordinate changes. A simple example of an invariant is the ratio of two distinct

distances of three points on a line [51]. Regardless of how the line is rotated, translated, and scaled, the ratio of any two distinct distances of the three points is constant. The application of invariants to the ATR problem is not new [24]. Invariants are constant functions by construction: if an object undergoes any allowable group transformation (or is expressed in an alternate coordinate system) invariant functions yield the same value. Invariant functions are expressible in an invariant basis that can be manipulated to remove undesirable variables, and they are specific to target properties (e.g., surface area and shape)[48].

For visible and infrared images, photometric invariants [28] and thermo-physical invariants [22] related to image properties (e.g., intensity and color) may be used to recognize objects. The SAR ATR problem relies almost exclusively on geometric invariants, i.e., invariants determined by the geometrical relationship between object points. These invariants are typically expressed using Euclidean distances, determinants, and inner products in the object space [38]. They are commonly ratios of determinants, where the numerators scale as surface areas and the denominators scale as volumes [48].

Any collection of N points that rotates and translates freely in a three-dimensional space has $3N - 6$ invariants. This expression is obtained because each point has three degrees of freedom, but three allowed rotations and three allowed translations eliminate six degrees of freedom. Thus, an object composed of four points P_1, P_2, P_3 and P_4 , where the point P_n is the vector $[x_n, y_n, z_n]^T$, has $3(4) - 6 = 6$ invariants. A useful standard transformation of the points that employs these six invariants, here designated $I_{21}, I_{31}, I_{32}, I_{41}, I_{42}$, and I_{43} , is

$$P_1' = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}; \quad P_2' = \begin{bmatrix} x_2 I_{21} \\ 0 \\ 0 \end{bmatrix}; \quad P_3' = \begin{bmatrix} x_3 I_{31} \\ y_3 I_{32} \\ 0 \end{bmatrix}; \quad P_4' = \begin{bmatrix} x_4 I_{41} \\ y_4 I_{42} \\ z_4 I_{43} \end{bmatrix}. \quad (1-1)$$

This transformation maps the first point to the origin, the second point to the x-axis, the third point to the x-y plane, and the fourth point to anywhere in the three-dimensional space. It is expressed as a standard 3 x 4 matrix formed by the invariants so that the indices on each invariant indicate its column and row in the matrix.

A statistical computer simulation quantifies the effects of noise and rotation on the performance of a simplified invariant classifier to justify the development of a more sophisticated classifier later in Chapter 4. A quadrilateral, which is a coplanar three-dimensional object, is chosen for simplicity, but it easily represents part of a target of interest (e.g., an airplane wing). The projection of each object consists of four random coplanar points in a 100 x 100 pixel image. The two invariant functions

$$\Phi_1 = \frac{\begin{vmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \\ z_1 & z_2 & z_3 \end{vmatrix}}{\begin{vmatrix} x_1 & x_2 & x_4 \\ y_1 & y_2 & y_4 \\ z_1 & z_2 & z_4 \end{vmatrix}}, \quad \Phi_2 = \frac{\begin{vmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \\ z_1 & z_2 & z_3 \end{vmatrix}}{\begin{vmatrix} x_1 & x_3 & x_4 \\ y_1 & y_3 & y_4 \\ z_1 & z_3 & z_4 \end{vmatrix}} \quad (1-2)$$

are formed; they are ratios of the determinants of three distinct points on the object and capture essential three-dimensional information [23].

In this statistical computer simulation, bivariate Gaussian noise with a standard deviation of one pixel in both the x and y directions is added to 18 rotations of each object point, in ten degree increments, to model the noise a typical video system may

experience in representing a three-dimensional object. Although it is simplistic to characterize all variations in the projected image as bivariate Gaussian noise in the image domain, the high variance of the added noise (which is 1% of the image size) should more than compensate to make this a pessimistic case. In this experiment, the invariants Φ_1 and Φ_2 identify all ten objects because the invariant point (Φ_1, Φ_2) of an object with noise is closer in Euclidean distance to the invariant point of the same object without noise than to the invariant point of any another object. In rare instances, when rotations cause the determinants of facets of interest to approach zero or cause degenerate objects (four nearly collinear points) to be created, the invariants change by as much as five percent from the invariants of the original noise-free objects, but all objects are still correctly identified.

When a four-point ($N = 4$) three-dimensional object is projected onto one dimension (a line), one-dimensional invariants may be formed, and translations along the line are the only transformations allowed. There are $N - 1 = 3$ invariant functions, here designated i_1, i_2, i_3 , that can each be expressed as differences in distance between any two distinct points.

Understanding the relationship between an object and its projection to a lower dimensional space gives further insight into the ATR problem [46]. If a High Resolution Radar (HRR) illuminates a three-dimensional object, the radar transforms the object points to two-dimensional image points. The radar transform may be eliminated and the object/image relationship between the three-dimensional invariants and the one-

dimensional invariants may be obtained as follows [48]. An unknown radar transform represented by the row vector $[\rho_1 \ \rho_2 \ \rho_3]$ creates the one-dimensional HRR profile

$$\begin{bmatrix} 0 & i_1 & i_2 & i_3 \end{bmatrix} = [\rho_1 \ \rho_2 \ \rho_3] \begin{bmatrix} 0 & I_{21} & I_{31} & I_{41} \\ 0 & 0 & I_{32} & I_{42} \\ 0 & 0 & 0 & I_{43} \end{bmatrix}, \quad (1-3)$$

where the radar transform is normalized so that

$$\rho_1^2 + \rho_2^2 + \rho_3^2 = 1. \quad (1-4)$$

Solving for the unknowns of the radar transform and substituting into Equation (1-4) yields

$$\begin{aligned} & [(I_{32}I_{41} - I_{31}I_{42})i_1 + I_{21}I_{42}i_2 - I_{21}I_{41}i_3]^2 \\ & + [(I_{31}i_1 - I_{21}i_2)^2 + (i_1^2 - I_{21}^2)I_{41}^2]I_{43} = 0. \end{aligned} \quad (1-5)$$

This equation provides an explicit relationship between the three-dimensional and one-dimensional invariants and shows that no knowledge of the radar transform is needed when exact knowledge of the object-image relationship is available, i.e., the equation relates 1D and 3D invariants, but it does not involve the radar transform elements.

1.1.2 Scattering centers

ATR algorithms commonly extract scattering centers from SAR images for data compression. The XPATCH program, developed elsewhere, employs a ray-tracing algorithm to solve Maxwell's equations at each boundary and determines a set of point scattering centers that can produce the given image when the target model is illuminated by a radar pulse from a predetermined aspect angle [60]. It is much more efficient to store the location and intensity of a number of point scattering centers than the entire

SAR image when building a database library of all possible targets, and this scattering center representation is commonly compressed even further by extracting a few useful features. However, ATR algorithms must input scattering centers (and often extract features) for many views of each target, since the position and intensity of the scattering centers vary significantly for small changes in aspect angle.

1.1.3 Existing two-dimensional classifiers

Some ATR algorithms rely on matching an unknown SAR image to templates of targets from a large database. Since a target is recognizable only within a few degrees of aspect angle, each target has many associated poses that are individually stored as templates. The unknown target is compared to templates within the library of possible targets and is matched to the closest template using a metric such as Mean Squared Error (MSE). A classifier using this method requires large amounts of memory to store the templates and requires significant time to compare them.

More sophisticated algorithms extract scattering centers from SAR images and derive useful features from the representation described in Section 1.1.2. Each SAR image is treated as a two-dimensional projection of the target, and the features of this projection are compared against a library of features of two-dimensional projections of known targets. Less memory is then needed for a classifier because the location and intensity of scattering centers are stored instead of the entire image, and the classifier is faster because it compares features of the images rather than entire images.

1.2 Three-dimensional MAGI concept

The 3D MAGI algorithm was not used in the research reported here, but it is discussed because it demonstrates the feasibility of extracting scattering centers which persist for up to 20 degrees of aspect angle and because it provides valuable insight into the design of a true three-dimensional classifier. In particular, the motion of the target relative to the sensor or the use of multiple sensors enables the extraction of the dominant three-dimensional scattering centers of the target. The SAR image is partially modeled by the extracted scattering centers; the part of the image not accounted for by the model is the residual. The representation of the target as a collection of points with known positions and intensities as a function of range explicitly determines three-dimensional invariants which describe the target independent of rotation, translation, and scale.

The target identification scenario proceeds from target data acquisition by SAR sensors to identification of the target using its three-dimensional invariant “fingerprint.” Multiple sensors are deployed for a stationary target to extract scattering centers from the raw SAR radar data (also known as the Video Phase History or VPH), and these scattering centers are tracked in a target-centered coordinate system to obtain a range history that details the relative movement of each scattering center on the target. These scattering center tracks are subtracted from the VPH to generate the residual VPH, and more scattering centers are successively extracted until an acceptable amount of the VPH is characterized. The multiple sensors generate multiple two-dimensional views, and the 3D MAGI algorithm provides three-dimensional estimates of the target scattering centers that persist over a wide range of aspect angles [48].

The classifier developed in Chapter 4 uses similar scattering center estimates to classify military targets. This classifier calculates invariant features from these estimates that describe the target three-dimensional properties independent of pose (angle of illumination), and this information is used to identify the target from a library of known targets. A new entry in the target library can be created for each new target, and if this new target is seen again, it can be recognized. Each successive viewing of a target generates more comprehensive three-dimensional information about the target.

1.2.1 Three-dimensional reconstruction

Any three-dimensional object of four points has six invariants (see Section 1.1.1). Any six views of this object create six equations in six unknowns. If these six views are non-degenerate, then they create six linearly independent equations that may be solved for the six invariants. If the sensor flies a marginally complicated flight path (e.g., it avoids tracing a conic section) or if the target moves in a similarly complicated path, the desired six non-zero eigenvalues in a system of linear equations is created. The relationship between the relative motion of the sensor and the number of non-zero eigenvalues of the linear system of equations is given in reference [48] and is presented in Table 1.

Figure 2 demonstrates how the three-dimensional model of a simple four-point object is generated from six views. In this example the four points of interest are the nose, back of the two wing tips, and tail (p_1 , p_2 , p_3 , p_4 , respectively) on an aircraft. To optimize a SAR flight path to recover all three-dimensional target information, rather than fly in a

straight line, the SAR sensor flies as complicate a flight path relative to the target as possible.

Table 1. Relative Motion vs Non-zero Eigenvalues of System

# of non-zero eigenvalues	Relative Motion of Sensor
1	Confined to a single line
2	Confined to two non-coplanar lines
3	Confined to a single plane
4	Confined to a pair of planes
5	Confined to an elliptic cone
6	Full 3-D motion

1.2.2 3D MAGI algorithm

After the raw SAR data are collected, many radar artifacts must be removed. This removal entails preparing a translation-compensated image that accounts for the relative motion of the target centroid from pulse to pulse. The goal is to provide a target-centered view over time.

Next, a number of scattering centers are extracted from the modified SAR data. Typically these scattering centers are prioritized by total energy, and the result is a range history of the extracted scattering centers over time. Additional scattering centers are extracted iteratively until the residual energy reaches an insignificant level.

The range history of the scattering centers commonly shows many scattering centers “crossing” over time. It is important to track each individual scattering center over time and establish a correspondence among all the scattering centers for the entire range history. The 3D MAGI algorithm uses velocity to track scattering centers when the range is ambiguous.




Figure 2. Three-dimensional reconstruction using six views. Here four points of interest are tracked in three dimensions and mapped to a useful standard position: the first point to the origin (P_1'), the second point to the x-axis (P_2'), the third point to the x-y plane (P_3'), and the fourth point to anywhere in the three-dimensional space (P_4'). Only these invariant mappings are guaranteed to yield the same result for any view of this object. Six linearly independent equations are solved to find the six three-dimensional invariants (I_{21} , I_{31} , I_{32} , I_{41} , I_{42} , I_{43}). Note that six non-degenerate views are required to find all six 3D invariants of a four-point object and capture its 3D nature, but only two views (represented as two curved lines) are shown in the figure.

Once the range histories for each scattering center are known, a three-dimensional scattering center model of the target may be reconstructed. The first step is a transformation from a space of dimension equal to the number of pulses to a space of dimension equal to the number of scattering centers. This step involves solving a least squares problem and using the Singular Value Decomposition (SVD) to find a number of largest eigenvalues equal to the number of scattering centers. The next step is to make a transformation to three-dimensional space with another SVD. Because the noise is uncorrelated in the higher dimensional space, a large number of scattering centers produces a better estimate of the object [48].

An efficient estimate of the scattering center coordinates in 3D may be generated: as the number of samples approaches infinity, the variance of the estimates approaches the population variance divided by the number of samples [49]. As new data is collected on a target (i.e., more pulses), the accuracy of the model increases, since every nondegenerate new view of the target provides more three-dimensional information.

In a demonstration conducted elsewhere, four corner reflectors were placed in a large flatbed truck such that they do not all lie in the same plane [48]. This arrangement formed the vertices of a tetrahedron, which illustrated the basic invariants discussed in Section 1.1.1. The truck moved in the footprint of a SAR platform to create the complicated motion necessary to extract three-dimensional target data. When given the raw SAR data of the target, the 3D MAGI Algorithm provided a near perfect reconstruction of the three-dimensional arrangement of corner reflectors.

For real targets the algorithm is set to extract enough scattering centers to effectively model the radar energy, since a real target has an unknown number of scattering centers and noise and background clutter can generate energy that may be modeled by low intensity scattering centers. Corner reflectors persist over a wide aspect angle range, and each of the four corner reflectors is visible in the experiment. However, real targets have some scattering centers that do not persist throughout a typical collection, and some real targets have axes of symmetry that generate multiple solutions.

1.3 Problem statement

The traditional deployment of SAR sensors involves flying a straight path over a stationary target. Although this geometry provides a clear two-dimensional projection of the target and sensor, a complicated relative motion of the target relative to the sensor yields three-dimensional information. In this case, the extracted scattering centers are three-dimensional (due to multiple noncoplanar projections), and the ATR algorithm may compare three-dimensional objects. A key goal of this dissertation is the design of a computationally efficient three-dimensional classifier to identify targets from SAR images and the development of an underlying theory for this design. This goal requires adopting or designing appropriate metrics to characterize three-dimensional targets.

1.3.1 Three-dimensional classification

The theory of classifying two-dimensional objects (e.g., images) is well understood [20]: objects are typically compared that are complete (i.e., every pixel has a value) and of the same size. However, after extracting three-dimensional scattering centers, the data form a sparse array (i.e., in a sampled three-dimensional space many voxels, the three-

dimensional equivalents of pixels, are empty). Each array can have a variable number of scattering centers, and the weight (intensity) of each of these scattering centers can vary. Finding useful statistics for these collections of scattering centers that enable ATR is a challenge.

Appropriate metrics to compare unknown targets to a library of known targets are needed before any three-dimensional classifier can be constructed. It is unclear whether traditional two-dimensional metrics such as MSE can be extended to this three-dimensional problem or whether new metrics must be developed. This dissertation determines the suitability of existing two-dimensional metrics and explores new techniques to compare three-dimensional targets.

1.3.2 Designing an optimal three-dimensional classifier

A classifier using extracted three-dimensional scattering centers and invariants is designed for eight military targets. The classifier should be optimal in that it uses features that are independent of any translation and rotation of the object in the radar image. Furthermore, this classifier should exhibit high data compression and should be able to identify both noisy and obscured targets. Since there is no precedent for a SAR classifier that uses true three-dimensional information, developing the theory of true three-dimensional SAR classification and implementing a system-level classifier design establishes the promise of this new technology. Four steps are necessary to implement a practical classifier: building scattering center models (Chapter 2), extracting features from the scattering centers (Chapter 3), populating the classifier training data with known targets (Chapter 4), and using the classifier to identify other known targets

(Chapter 5).

After constructing a true three-dimensional SAR classifier, its performance is compared to that of an existing two-dimensional classifier. The memory requirements and speed of each classifier are determined (using a common computer for both classifiers to compare an unknown target to an existing target in the database). The two types of classifiers are compared for “real world” conditions, specifically white noise and target obscuration. The three-dimensional classifier is expected to yield significant improvement over the two-dimensional classifier, since a two-dimensional classifier must store multiple projections of three-dimensional targets (or use processor time to reconstruct them).

2 *Model Building*

The 3D MAGI algorithm is computationally intensive and is not compatible with available ray-tracing algorithms. More efficient procedures for creating discrete scattering center models of known targets are developed in this chapter for use with the invariant-feature classifier developed and discussed in the following chapters.

Any complex radar return (i.e., a return that has both amplitude and phase components) can be modeled as a sum of scattering centers (see Section 1.1.2) that function as ideal corner reflectors [39]. If the coordinates of these scattering centers cannot be easily extracted from a complex radar image, then it may not be feasible to develop a good classifier due to the difficulty of populating a database of target models.

An effective and tractable model for two-dimensional scattering has been developed elsewhere in which the sensor moves in a line perpendicular to the target sensor direction [13]. Here this model is extended to three-dimensional scattering in which the sensor moves nonlinearly in a closed path over a few steradians in a plane perpendicular to the target sensor direction, the target is stationary, and the sensor is in the far field (i.e., many wavelengths distant). These assumptions are consistent with data collection from a side-looking SAR platform and are easily satisfied in a controlled experiment.

2.1 *Three-dimensional scattering model*

A radar return or pulse is modeled as the sum of returns from N scattering centers. The amplitude A_n and coordinates (x_n, y_n, z_n) of each scattering center are unknown. The amplitudes do not significantly change over small angles, and thus they are considered constant from pulse to pulse. The coordinate system imposed on target and sensor has its

origin on the target center and has its x-axis in the sensor direction for the first measurement. Also, azimuth is in the ϕ direction (counterclockwise is positive), and elevation is in the θ direction (upwards is positive). The electric field strength E_n of each individual scattering center and the total electric field strength E_{TOTAL} are then extended to three dimensions:

$$E_n = A_n j \frac{f}{f_c} \exp\{ -2\pi f \gamma'_n \sin(\phi - \phi'_n) \gamma''_n \sin(\theta - \theta'_n) \} \exp\{ j \frac{4\pi f}{c} (x_n \cos \phi \cos \theta - y_n \sin \phi - z_n \sin \theta) \} \quad (2-1)$$

$$E_{TOTAL} \cong \sum_{n=1}^N A_n j \frac{f}{f_c} \exp\{ j \frac{4\pi f}{c} (x_n \cos \phi \cos \theta - y_n \sin \phi - z_n \sin \theta) \}. \quad (2-2)$$

Here A_n is the electric field amplitude of center n , f is frequency of the swept pulse, f_c is its center frequency, γ'_n is the amplitude damping coefficient in the ϕ direction, and γ''_n is the amplitude damping coefficient in the θ direction, where the damping coefficients describe the attenuation of the electric field amplitudes in the ϕ and θ direction.

Invariant features of the resulting discrete set of three-dimensional scattering centers can be used to classify targets. These invariant features are, by definition, independent of scattering center translation, rotation, and scale. A different choice of coordinate axes in the collection geometry introduces a rotation matrix that does not affect invariant features. Also, fixing one scattering center at the origin and calculating relative distances to the other scattering centers imposes an arbitrary translation that does not change invariant features.

2.2 Extraction procedure

An effective model must represent a realistic radar system in an actual collection mode. Here the radar system is assumed to operate at 8 - 12 GHz and at over 10 nautical miles (18.4 km) from the target. It also assumed that all targets of interest are no larger than 30 m in any direction and that at any time the exact orientation of the target relative to the sensor is known. The latter assumption is valid because target models are generated under controlled conditions (e.g., a radar turntable experiment). Exact knowledge of relative sensor position permits more efficient scattering center extraction than the 3D MAGI algorithm.

2.2.1 Extracting amplitude and “wrapped” phase

The first step is to determine the peak amplitude of each scattering center and its associated phase from the Fourier transform of the radar return. The exact position on the viewing sphere, the frequency of the radar, and the total returned signal are known. The magnitude and phase of each vector that sum to the total electromagnetic field, E_{TOTAL} in Equation (2-2) are unknown.

A Fourier transform of a sum of complex exponentials in the spatial domain yields a sum of impulses in the frequency domain. Phase shifts in the spatial domain correspond to translations in the frequency domain. The peak amplitude of each scattering center is extracted from the Fourier transform of the radar return, where the relative intensity of each scattering center is preserved. Since the exponentials are modulated on a carrier, the carrier is removed to bring the spectrum back to baseband.

The phase difference between scattering centers is determined from the difference between the impulse positions in the frequency domain, which is easier than extracting the phases of many complex exponentials summed in the spatial domain. However, the phase difference between scattering centers is “wrapped” in that it must lie in the interval $(-\pi, \pi]$, i.e., from greater than minus pi to less than or equal to π . The “unwrapped” phase necessary to solve for the location of scattering centers differs from the wrapped phase by a multiple of 2π for each complete half wavelength of separation along the viewing angle.

2.2.2 Unwrapping phase

The absolute phase difference of any two scattering centers from one sample on the viewing sphere to the next must differ by less than π radians; otherwise this difference is ambiguous because it repeats every 2π radians. If the maximum extent of the target l_{\max} in multiples of half the radar wavelength is known, the minimum sampling distance on the viewing sphere is easily obtained as

$$\Delta \phi_{\min} = \arccos\left[\frac{(2l_{\max} / \lambda) - 1}{(2l_{\max} / \lambda)} \right]. \quad (2-3)$$

The worst case occurs when the distance from one scattering center to another is the length of the target and the viewing angle is parallel to the distance vector from one scattering center to the other. In this case, the maximum allowed difference in the length of the two paths is one half wavelength. As shown in Figure 3, the hypotenuse of the right triangle is the extent of the target, and the longest leg is exactly one half wavelength less. Thus the argument to be minimized is an inverse cosine, and the worst case is when

this argument is as large as possible. For the assumptions made in Section 2.2, this case corresponds to a maximum sampling interval of 0.0326° .

For sufficiently small sampling intervals, “wrapped” phase differences that are modulo 2π are converted to “unwrapped” phases. All phase differences are constrained to lie in the interval $(-\pi, \pi]$, so a phase difference of $3\pi/2$ is converted to $-\pi/2$. The first sample is chosen as an arbitrary reference for which the “unwrapped” phase is set equal to its “wrapped” phase. The “unwrapped” phase for sample two is the “wrapped” phase for sample two plus the phase difference between “wrapped” sample two and “wrapped” sample one. The “unwrapped” phase for sample three is the “wrapped” phase for sample two (just calculated) plus the phase difference between “wrapped” sample three and “wrapped” sample two. This procedure continues until all the phase terms are “unwrapped” over the entire sampling angle.

2.2.3 Solving for scattering center coordinates

Multiple phase measurements and the associated angular positions can be represented as an “unwrapped” phase vector ($\Phi_{\text{unwrapped}}$) at a vector of known positions (ϕ). With “unwrapped” phase known, these vectors are related to the relative scattering center coordinates Δx and Δy by

$$\Phi_{\text{unwrapped}} = \Delta x \cos \phi - \Delta y \sin \phi . \quad (2-4)$$




Figure 3. Two difficult cases for “unwrapping” phase. Here $\Delta\phi_{\min}$ is the minimum sampling interval needed to unambiguously “unwrap” the phase between two scattering centers, where the total phase change must be no greater than one half wavelength. The maximum length of the target forms the hypotenuse of a right triangle, and the maximum length of the target minus one forms the long leg of this triangle, where all lengths are expressed in half wavelengths. The expression to be minimized is an inverse cosine, and the worst case is when the distance between scattering centers is a maximum and the position vector from one scattering center to the other is parallel to the look angle. The diagram on the right shows the only case where the phase difference between two scattering centers changes sign from one sample to the next.

Solving (2-4) for the unknown vector $[\Delta x \ \Delta y]^T$ yields the overdetermined system

$$\begin{bmatrix} \Phi_1 \\ \Phi_2 \\ * \\ * \\ * \\ \Phi_N \end{bmatrix} = \begin{bmatrix} \cos \phi_1 & -\sin \phi_1 \\ \cos \phi_2 & -\sin \phi_2 \\ & * \\ & * \\ & * \\ \cos \phi_N & -\sin \phi_N \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}. \quad (2-5)$$

If A is the full rank matrix of sines and cosines above (the nonlinear behavior of the sine and cosine should ensure A is full rank), then the least squares solution is obtained by multiplying both sides of Equation (2-5) by A^T from the left and solving the resulting two equations in two unknowns for $[\Delta x \ \Delta y]^T$. Also note that for the classifier designed here or any classifier that uses geometric features, the phase terms (and the resulting coordinates) are sufficient, and the amplitude of each scattering center is not characterized.

The above algorithm determines the separation between the reference scattering center (the highest intensity one) already mapped to the origin and the next highest intensity scattering center. The original spatial frequency spectrum is used to determine the separation between the reference scattering center and each successive scattering center until a sufficient amount of the radar energy is modeled.

Sampling along the ϕ direction determines the scattering center coordinates in the x-y plane, but rotation around the z-axis does not resolve separation in the z dimension. Thus, the above steps are repeated while sampling in the θ direction to determine the Δx and Δz coordinates. The Δx coordinates are already calculated, so consistency may be verified.

2.3 *Algorithm validation and test*

To validate the extraction algorithm, a five-point object in three dimensions is generated from a uniform distribution in a 30-meter cube consistent with the assumptions of Section 2.2. The complex radar return is calculated from Equations (2-1) and (2-2) using eleven samples of ϕ from 0° to 1° in 0.1° increments and using unit amplitudes A_n

and zero damping coefficients γ' and γ'' . The phase differences “wrapped” between one scattering center and the other four scattering centers are extracted from the spatial frequency spectrum by calculating distances from one impulse to the next.

Equation (2-3) gives the minimum sampling interval for this test as 0.0326° .

Although the data is not sampled finely enough, some properties of the collection geometry may be used to obtain better results. Figure 4 shows that if the collection geometry does not include a broadside case between any two scattering centers, the “unwrapped” phase must be monotonically increasing or decreasing. Thus, there are two competing hypotheses: the phase differences lie in either the interval from 0 up to 2π or the interval from 0 down to -2π . After performing the least squares fit to solve Equation (2-5) using both assumptions, only the valid hypothesis fits the model (see Figure 4).

In the test it is not possible to unambiguously unwrap phase due to the overly large sampling interval on the viewing sphere. However, by testing the hypotheses that the phase is monotonically decreasing or increasing, it is possible to reject the incorrect hypothesis, to recover the separation of the scattering centers, and to correctly unwrap the phase. The $[\Delta x \ \Delta y]^T$ has the lowest mean squared error between the hypothesized phase vector and that calculated in Equation (2-5). This technique effectively increases the minimum sampling interval by a factor of two, but fails to correctly unwrap phase when one of the samples is broadside to two scattering centers (see Figure 3).

To account for this rare possibility, a number of hypotheses equal to twice the number of samples may be tested to determine the scattering center separation by locating the sample at which phase difference changes sign. This procedure requires solving for a




Figure 4. Unwrapping phase between two scattering centers on a larger object. A random object is constrained to lie in a 30-meter cube, which corresponds to a 2,000 x 2,000 x 2,000 half wavelength space. The graph shows the “unwrapped” phase of two possible solutions given monotonically increasing or monotonically decreasing phase terms. Only one solution models the actual data and recovers the separation between the two chosen scattering centers.

number of $[\Delta x \ \Delta y]^T$ vectors equal to two times the number of samples and assuming that there is exactly one sign change in the phase difference vector. Once again the correct $[\Delta x \ \Delta y]^T$ vector is the one with lowest mean squared error between the hypothesized phase vector and that calculated from a least squares solution of Equation (2-5). Thus, using the algorithm detailed in Section 2.2, the scattering center coordinates of a random object are completely extracted, even though the sampling interval is over three times the minimum sampling interval needed to guarantee unambiguous phase “unwrapping”. No

broadside case exists in the collection geometry, so only two competing hypotheses require testing for each scattering center coordinate.

2.4 Implementation

Here three-dimensional reconstruction has been treated as two two-dimensional reconstructions. This treatment requires only one-dimensional (instead of two-dimensional) Fourier transforms. Experiments to validate the algorithm used N samples on the viewing sphere in one direction to solve for scattering center coordinates [17]. A three-dimensional reconstruction at the equivalent sampling interval would use N^2 samples on the viewing sphere to obtain a much higher ratio of equations (samples) to unknowns, and it would not solve twice for the same coordinate (Δx in the sample collection geometry used here).

An ambitious plan to apply this algorithm (later used in Section 4.3) would build one large target model valid over all possible angles on the viewing sphere. Scattering centers would have associated two-dimensional rectangular window functions that would model the persistence of each scattering center. At any point on the viewing sphere a set of scattering center coordinates describing the object could be easily obtained. Any unknown target sampled over a subset of the total viewing sphere could be compared to existing models of valid targets, and the best match in three-dimensional shape could be determined (ideally with associated confidence measures).

2.5 *Impact on classifier design*

Given complete knowledge of any target selected for a three-dimensional classifier, there is now an effective and efficient way to populate a target database with target models to be implemented in Chapter 4.

Commercial ray-tracing algorithms such as XPATCH can calculate the complex radar return from a CAD model of a known target [60]. After the CAD models are imported to the ray-tracing algorithm, the location of a radar transmitter (in aspect angle, elevation angle, and distance) is specified along with the parameters of the radar (center frequency, number of frequency steps, etc.). The algorithm solves Maxwell's equations at object boundaries to determine the energy collected at a receiver at a specified location (in aspect angle, elevation angle, and distance). Ray-tracing algorithms often have the ability to extract the scattering center representation of an object that generates the radar image (commonly represented as a Video Phase History), but they do not extract the most stable scattering centers as a function of aspect angle, only the high-intensity ones. Thus these scattering center representations are far from robust, and break down trying to model the simplest of targets (e.g., a few corner reflectors).

If such ray-tracing algorithms are optimized to extract rigid three-dimensional scattering centers, a complete database of models can be built entirely on high-speed computers. Alternatively, inexpensive small-scale models of known targets can be built and placed on a compact radar range (using miniaturized targets and reduced radar wavelengths). The complex radar returns from an experiment where the small-scale model is rotated and possibly tilted would allow the extraction algorithm to build a model

for a database of known targets. Either of these two methods constructs “cloud of points” target models efficiently and justifies the design of a three-dimensional classifier.

3 Singular Value Features

Appendix 2 examines using Procrustes distance as a metric for classifying targets modeled by sets of scattering centers. Although Procrustes distance is invariant to translation, rotation, and scale, matching objects with unequal numbers of scattering centers is problematic. This chapter considers singular value features, which avoid this matching problem but have the desired invariance properties.

3.1 Definition of singular values

Eigenvectors and eigenvalues that describe coordinate rotations and scalings, respectively, can be obtained for square matrices. For non-square matrices Singular Value Decomposition (SVD) may be employed for these descriptions. Given an $m \times n$ matrix A with $m \geq n$, SVD extracts unitary matrices $U_{m \times m}$ and $V_{n \times n}$ and diagonal matrix $S_{m \times n}$ (at least the top n rows of $S_{m \times n}$ form a diagonal matrix) such that $A^H V = U^H S$, where the top n rows of $S_{m \times n}$ have the n singular values of matrix A on its diagonal (most SVD algorithms order the singular values from largest to smallest) and where H designates complex conjugate transpose. If matrix A represents $m \geq n$ points in an n -dimensional space, the n singular values are the lengths of the n semi-axes of an n -dimensional ellipsoid. This ellipsoid defines the one standard deviation contour of a zero-mean Gaussian probability density of covariance matrix A . This covariance matrix is $m^{-1} A_0^H A_0$, where A_0 is A with the mean of each column subtracted from the elements of that column and where the singular values are the square roots of the eigenvalues of the covariance matrix.

3.2 Invariance of singular values

Any translation or rotation of the set of points represented by matrix A does not change the length of the semi-axes of the n -dimensional ellipsoid discussed previously. However, any scaling (i.e., multiplication of each coordinate by the same constant) of these points scales the n lengths of the semi-axes by the same constant. Thus, singular value features are invariant to translation and rotation of the original matrix but preserve scaling.

Given a three-dimensional object matrix $A_{n \times 3}$, the three singular values of the matrix roughly correspond to half the length, half the width, and half the height of the object. These three singular values may therefore be suitable features for target identification.

3.3 Monomial expansion of object matrix

Methods for obtaining higher order expansions of the object matrix are desired, since the singular values of such expansions define the object matrix more completely than the original three singular values. In particular, a large number of singular values contains the same information as the original x , y , and z values of the original object matrix.

A monomial expansion is similar to a binomial expansion, except the numerical coefficients of the terms are unimportant. If each row (i.e., each scattering center point) of an $N \times 6$ object matrix is represented as $[x_n \ y_n \ z_n]$, then each row of the second order monomial expansion of the object matrix is $[x_n^2 \ y_n^2 \ z_n^2 \ x_n y_n \ x_n z_n \ y_n z_n]$. If the coordinates of the object matrix are converted to zero mean, the first and last three columns of this second order monomial expansion are related to coordinate variances and

covariances, respectively. The six singular values of this $N \times 6$ matrix provide additional information about the object.

Further monomial expansions of the object matrix improve the representation. Just as skew and kurtosis provide more information (related to the third and fourth moments of a probability distribution), further expansions of the object matrix provide more information about the original object. However, if each coordinate consists of a deterministic value and a noise component, the higher the order of the monomial expansion, the higher the noise content of each element in the expanded matrix. In the following, a simulation shows that the 19 singular values of the first, second, and third monomial expansions of the object matrix are sufficient to achieve adequate classification performance.

3.4 Determining the feature set

To select a suitable feature set, a group of ten simple ten-point objects is created from three coordinates distributed independently and uniformly from zero to one meters, and simple Euclidean classifiers are tested based on the singular values of monomial expansions of increasing order. For example, the simplest feature set is the three-dimensional vector composed of the three singular values of each object matrix, and a simple Euclidean distance classifier matches the three singular values of any unknown object to the closest singular values of a known object. The next higher-order feature set is the nine-dimensional vector consisting of the last feature set augmented by the six singular values of the second order monomial expansion of this matrix.

All singular value classifiers are compared to a benchmark classifier that treats every coordinate in the object matrix as a feature. This thirty-dimensional feature set contains an exact representation of each object. As successive classifiers add more singular values, classifier performance improves, but there is a limit where added features contribute little to classifier performance.

Since there is only one example of each object, each classifier is tested using a Monte Carlo method in which equal numbers of each of the original ten objects have white noise added to the ten points. Noise variance is successively increased to degrade classifier performance, and any unknown (noise-corrupted) target is matched to the known (noise-free) target closest in Euclidean distance in the feature space. Under ideal conditions (no noise), the classifier correctly classifies every target (since without noise there are no unknown targets). For high noise levels, classification performance tends to 10% (that of a random guess). Figure 5 shows how each classifier compares to the benchmark classifier that uses the coordinates of each of the ten object points as features. For the ten-point objects, the correct-classification performance of a classifier that uses only the first three singular values is about 60% that of the benchmark classifier for a noise standard deviation of 0.5m. Using six singular values yields about 90% of the performance of the benchmark classifier for this noise standard deviation, and using ten singular values improves performance to over 98% that of the benchmark classifier. Thus, for this example, augmenting the feature space by generating further higher order monomial expansions is unnecessary.




Figure 5. The significance of features from successive monomial expansions. Here ten three-dimensional objects, each composed of ten points, are created such that each coordinate is uniformly distributed between zero and one meter. The objects are then subjected to noise, where the noise standard deviation is that of white Gaussian noise added independently to each coordinate. Four nearest-Euclidean-distance classifiers of different complexity are tested: a classifier that uses all elements of the ten by three object matrix, a classifier that uses only the three singular values of each ten by three object matrix, a classifier that uses the three singular values from the ten by three object matrix plus the six singular values of its second order monomial expansion, and a classifier that uses the three singular values of the object matrix and both the six singular values of its second order monomial expansions and the ten singular values of the third order monomial expansion. The latter classifier has a 97% probability of correctly identifying an unknown target at an intermediate noise level (0.5m standard deviation) compared to 98% for the benchmark classifier that uses all information about the object (i.e., the coordinates of the ten points). Thus further monomial expansions are unnecessary.

Similar tests that used a view of each of eight military objects with additive zero-mean white Gaussian noise produced similar results. In particular, the probability of correct classification appreciably increases until the feature set includes the singular values of the target matrix and its second and third order monomial expansions. However, adding the singular values of the fourth order monomial expansion significantly increases the dimensionality of the feature space with only a slight improvement in classifier performance.

3.5 *Feature set definition*

Section 3.4 indicates that a nineteen-dimensional vector calculated from the singular values of the target matrix and its first two monomial expansions (the second and third order expansions) should contain adequate information for effective classification. For complicated objects (objects with large numbers of scattering center points), the performance of classifiers that use this feature vector should decrease. Anecdotal evidence suggests that certain scattering centers are more persistent as a function of aspect angle than others [53]. Thus, if the ten scattering centers that persist for the widest range of aspect angles are extracted from a much larger set, and if the remaining scattering centers are not nearly as persistent, then the performance of a classifier using the nineteen-dimensional vector should be close to that of a classifier using all coordinates.

In this dissertation, the nineteen-dimensional feature vector is chosen to represent diverse military targets (see Figure 6). The elements of the feature vector are the singular values of the object matrix augmented with the singular values of its second and third

order monomial expansions and ordered from largest to smallest. This feature vector is independent of translation and rotation of the original object matrix. If the object matrix is scaled by k , the elements of the feature vector are scaled by k , k^2 , or k^3 depending on their associated monomial expansion. Thus, this feature vector preserves scaling and can discriminate between objects of different physical size. Because the size of the target in the SAR image is independent of the distance from the target to the SAR sensor, this preservation of size is advantageous (which would not be the case for an optical system where size and distance to the target are related and invariance to scale may be desired) [50].

Concatenating the singular values of three matrices forms the selected feature vector, but concatenating the three monomial expansion matrices first and then finding the singular values of this single nineteen-column matrix forms an alternative feature vector. This feature vector is not invariant to translation and rotation of the original object matrix and is not considered here.

The chosen feature vector is not only independent of translation and rotation of the object matrix, but also, by construction, guarantees perfect correspondence between feature vectors from object matrices with different numbers of points. Perfect correspondence means that every object can be described by exactly nineteen singular values regardless of the number of scattering center points (typically thousands).

4 Classifier Design

The three-dimensional scattering centers of a target generated from raw SAR data can be represented in a nineteen-dimensional feature space. The development of an effective algorithm to classify targets in this nineteen-dimensional feature space is explored in this chapter. Understanding how this classifier functions is important because scattering center extraction has heretofore emphasized image compression, which may not be appropriate for SAR target classification.

4.1 Classifier design and compromises

The problem addressed here requires the correct identification of a set of eight targets. An ideal target might consist of a collection of scattering centers that exhibit constant reflection at all aspect angles and that are of sufficient number to provide a unique and complete object description. However, here eight challenging military targets and their associated scattering centers form the target set for a classifier that address a “real world” problem.

Existing classifiers treat each view of a target as an independent observation. Each SAR image has a time series of associated aspect angles, and regarding this series as the result of a known rotation makes use of a priori information that improves classifier performance.

4.2 Scattering center stability

Current ray-tracing algorithms are not capable of extracting scattering centers that are stable beyond three degrees, where a scattering center (in a target-centered frame of

reference) is stable if it is completely visible over a specified range of aspect angles. However, with the 3D MAGI algorithm described in Section 1.2, scattering centers from SAR returns that are stable over twenty degrees of aspect angle may be extracted. These scattering centers behave as corner reflectors with a two-dimensional damped sinusoid amplitude profile [13] described by Equation (2-1). It is not necessary to know the five parameters for each scattering center, i.e., maximum amplitude (A_n), location of the maximum amplitude in phi and theta (θ'_n, ϕ'_n), and damping rates in phi and theta (γ'_n, γ''_n). Extracting only the three-dimensional location (x_n, y_n, z_n) of each scattering center of the target is necessary.

The stability of the scattering centers is more important than their intensity (proportional to energy) when extracting scattering centers from SAR images of real objects and CAD models. For the singular value features used here (and any other geometric features), the intensity of the scattering centers is ignored.

4.3 Target models

Here eight vehicle targets of military interest (see Figure 6), the Russian SA-8 surface-to-air missile battery, the Russian SA-13 surface-to-air missile battery, the Russian T-72 main battle tank, the Russian T-62 tank, the US M-1 Abrams main battle tank, the US M-60 tank, the US M-113 armored personnel carrier, and the US M-911 ten-ton truck, are input to XPATCH software, which is commonly used ray-tracing algorithm that calculates scattering center locations [60]. However, each scattering center is constrained to function as a corner reflector modeled by Equation (2-1). Although the model is valid for targets over the entire viewing sphere, the elevation angle is assumed



Figure 6. Eight military vehicles used as targets in the singular value monomial expansion classifier. They are from left to right and top to bottom: the Russian SA-8 surface-to-air missile battery, the Russian SA-13 surface-to-air missile battery, the Russian T-72 main battle tank, the Russian T-62 tank, the US M-1 Abrams main battle tank, the US M-60 tank, the US M-113 armored personnel carrier, and the US M911 ten-ton truck.

fixed at just above zero degrees, since the data are for ranges of more than 10 nautical miles.

Exhaustive XPATCH runs generate discrete target representations every tenth of a degree in azimuth, where all scattering centers are constrained to persist for twenty degrees of azimuth angle. This modification to XPATCH allows for anticipated improvements in scattering center models that future algorithms such as 3D MAGI are expected to provide. A single four-column matrix describes each target: the first three columns are three-dimensional points that indicate all visible scattering centers, and the last column is the azimuth angle of the center of each scattering center “window”. Any scattering center at an azimuth angle within ten degrees of the center of a scattering center “window” is visible.

4.4 Singular value feature extraction

Each view of a target is modeled as a set of scattering centers, and features for classifier design are selected from these sets. As discussed in Section 3.3, a monomial expansion of the target matrix generates higher order moments (see Figure 7). If the singular values of these expanded matrices are used to augment the feature space, additional information about the targets is obtained. As indicated in Section 3.4 for simple ten-point objects, the singular values of the target matrix and its first two monomial expansions form a nineteen-dimensional feature vector that functions almost as well as the scattering center coordinates. For real targets there may be one thousand or more scattering centers.

$$\begin{bmatrix} x_1 & y_1 & z_1 \\ x_2 & y_2 & z_2 \\ \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots \\ x_N & y_N & z_N \end{bmatrix} \quad \begin{bmatrix} x_1^2 & y_1^2 & z_1^2 & x_1 y_1 & x_1 z_1 & y_1 z_1 \\ x_2^2 & y_2^2 & z_2^2 & x_2 y_2 & x_2 z_2 & y_2 z_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_N^2 & y_N^2 & z_N^2 & x_N y_N & x_N z_N & y_N z_N \end{bmatrix}$$

Target Matrix Second Order Monomial
(Nx3) Expansion of Target Matrix (Nx6)

$$\begin{bmatrix} x_1^3 & y_1^3 & z_1^3 & x_1^2 y_1 & x_1^2 z_1 & \cdot & \cdot & \cdot & x_1 y_1 z_1 \\ x_2^3 & y_2^3 & z_2^3 & x_2^2 y_2 & x_2^2 z_2 & \cdot & \cdot & \cdot & x_2 y_2 z_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_N^3 & y_N^3 & z_N^3 & x_N^2 y_N & x_N^2 z_N & \cdot & \cdot & \cdot & x_N y_N z_N \end{bmatrix}$$

Third Order Monomial Expansion
of Target Matrix (Nx10)

Figure 7. Monomial expansions of the target matrix. There are three singular values from the target matrix, an additional six singular values from the target matrix augmented with its second order monomial expansion, and an additional ten singular values from the target matrix augmented with its second and third order monomial expansions. The nineteen singular values of all these matrices form a nineteen-dimensional feature space.

4.5 *Creating target subclasses*

Each target is divided into subclasses in the feature space. The closeness of two aspect angles is related to the correlation of the associated sets of scattering centers. Since all scattering centers persist over twenty degrees of aspect angle by construction, two aspect-angle views separated by twenty degrees are completely uncorrelated. However, any aspect-angle view between these two views must consist only of scattering centers from these two views. A view exactly halfway between these views is expected to contain half of the scattering centers from one view and half of the scattering centers from the other view. Also, for any arbitrary aspect angle between the two views, the expected number of scattering centers from each of the two views can be calculated: the distance from the arbitrary aspect angle to each view divided by twenty degrees is approximately the fraction of scattering centers that should be visible (see Figure 8).

Thus the 360° range of aspect angle is separated into 36 subclasses, each of which represent models generated over ten degrees of aspect angle. This separation ensures that the correlation between the scattering centers of any two aspect-angle views in the same subclass is at least 50% (see Figure 9). The 360° range of aspect angles is sampled every tenth of a degree to generate a population from which to calculate results.

4.6 *Classifier training*

There are 100 target models of each of 36 subclasses, and 5% of these models are withheld for later classifier testing. The nineteen-dimensional feature vector for each of




Figure 8. Plot of visible scattering centers as a function of aspect angle. Each target has 1234 visible scattering centers at zero degrees and 1144 visible scattering centers at twenty degrees. Note that the number of visible scattering centers as a function of aspect angle has the expected linear trend.




Figure 9. Correlation as a function of aspect angle. The graph shows the correlation between any two sets of aspect angle views of the same target as a function of difference in aspect angle. Since views separated by more than twenty degrees are uncorrelated, the 360° range of aspect angles may be separated into 10° subclasses, in which case the correlation between any two aspect angles is no greater than 0.50.

the 95 remaining models is calculated. From these 95 vectors, a mean vector and a covariance matrix are determined for each subclass of each target.

If the feature vector of each subclass is normally distributed, its distribution is completely described by a length nineteen mean vector and a nineteen by nineteen covariance matrix. The improved stability (persistence) of the scattering centers (up to 20 degrees) enables a higher correlation within subclasses that has been previously unachievable, and thus each subclass is reduced to a mean vector and a covariance matrix as shown in Figure 10. The resulting data compression is remarkable: each subclass that originally consisted of one hundred radar images (each about 12 megabytes) is now represented by only 209 numbers (19 mean vector elements plus 190 independent elements in the covariance matrix) that take about 800 bytes to store as double precision variables.

4.7 Decision Making

After characterizing all known targets as nineteen-dimensional Gaussians, a decision rule for classifying any unknown set of scattering centers is needed. The unknown set must be reduced to a nineteen-dimensional vector by taking the singular values of the target matrix and its monomial expansions. Rather than classifying this unknown vector \mathbf{x} by finding the closest subclass mean $\boldsymbol{\mu}$ in Euclidean distance, it is classified by finding the closest subclass mean in Mahalanobis distance r . The square of this distance is

$$r^2 = (\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}), \quad (4-1)$$

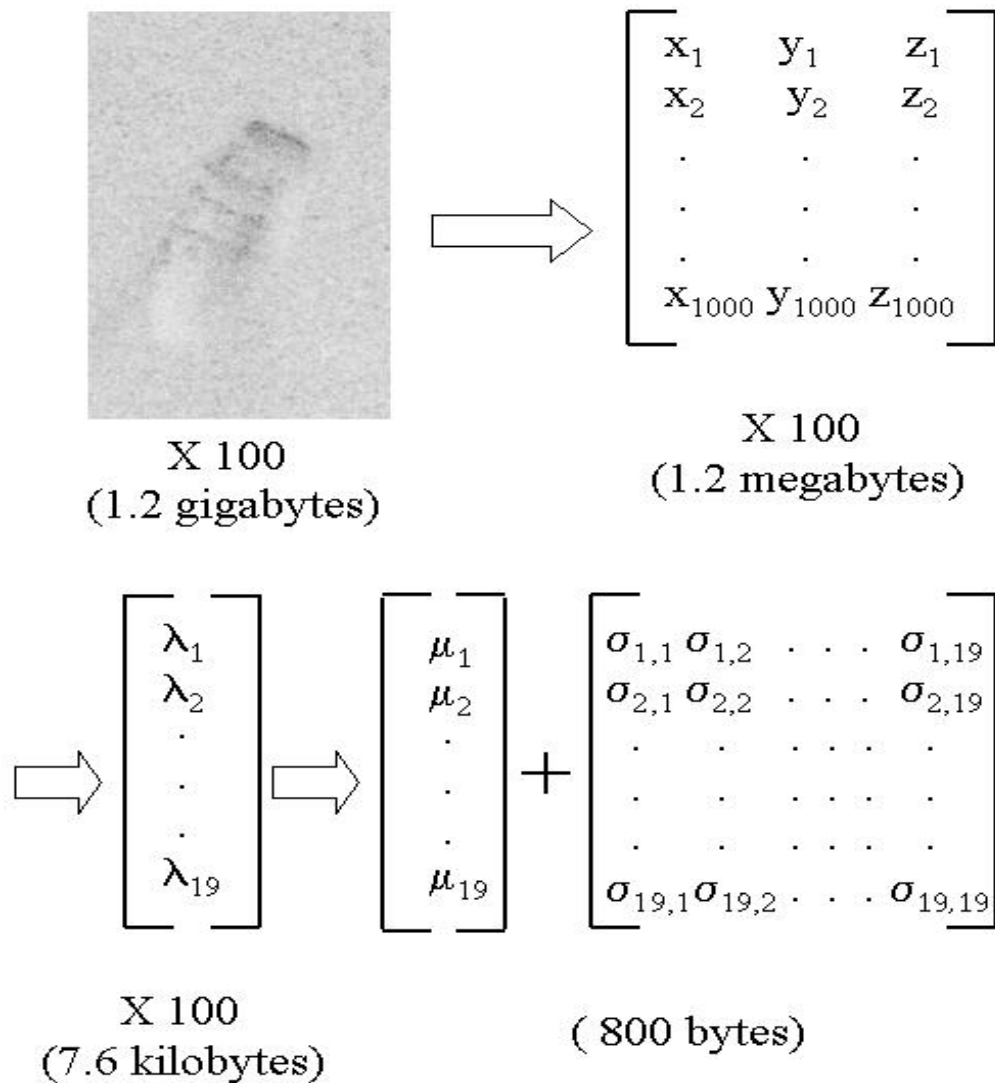


Figure 10. Data compression illustration. Each subclass originates from one hundred radar images (one image every tenth-degree in azimuth, with each subclasses encompassing ten degrees). A scattering center extraction algorithm that finds 1000 scattering centers reduces this data to one hundred 1000 by three matrices. A nineteen dimensional vector of singular values in the feature space represents each of these matrices, and finally, all one hundred nineteen-dimensional vectors are represented as a nineteen-dimension mean vector and a symmetric nineteen by nineteen covariance matrix. Thus, about one gigabyte of data is reduced to about 800 bytes. (Here the SAR image is not of a military target, but a car at a high elevation angle.)

where Σ is the class covariance matrix. Thus r^2 is squared Euclidean distance weighted by the inverse of the covariance matrix, where this matrix is normalized (divided by its determinant) to avoid scaling problems. Mahalanobis distance is used because it measures the number of standard deviations in the desired direction to account for the hyper-ellipsoid profile of the probability density associated with each class.

The decision rule takes any unknown feature vector and classifies it as one of the eight known targets. However, this procedure does not reject clutter or targets that are not in the database, so another rule to reject false targets is required. For each subclass, the Mahalanobis distance to all subclasses for a different target is calculated. The smallest of these distances is an estimate of a threshold that, when exceeded, indicates that the unknown vector is likely not associated with any of the known targets in the database. The decision process is shown as a flowchart in Figure 11.

There is a single conditional decision rule: an unknown vector is matched to the target of the associated subclass nearest to the unknown in Mahalanobis distance, provided this distance does not exceed the distance from the associated subclass to the nearest subclass of any different target, in which case it is a non-target (e.g., clutter). This proposed clutter threshold is a worst-case limit; Section 7.2.3 discusses improved decision-making procedures. Figure 12 is two-dimensional illustration in which the unknown is a single standard deviation from Target 1 View 3, four standard deviations from Target 1 View 14, and six standard deviations from Target 4 View 7. Therefore, the unknown is matched to Target 1 View 3. The clutter threshold is the distance to the closest view of a

different target: approximately four standard deviations in this example (the distance from the mean of Target 1 View 3 to the mean of Target 4 View7).

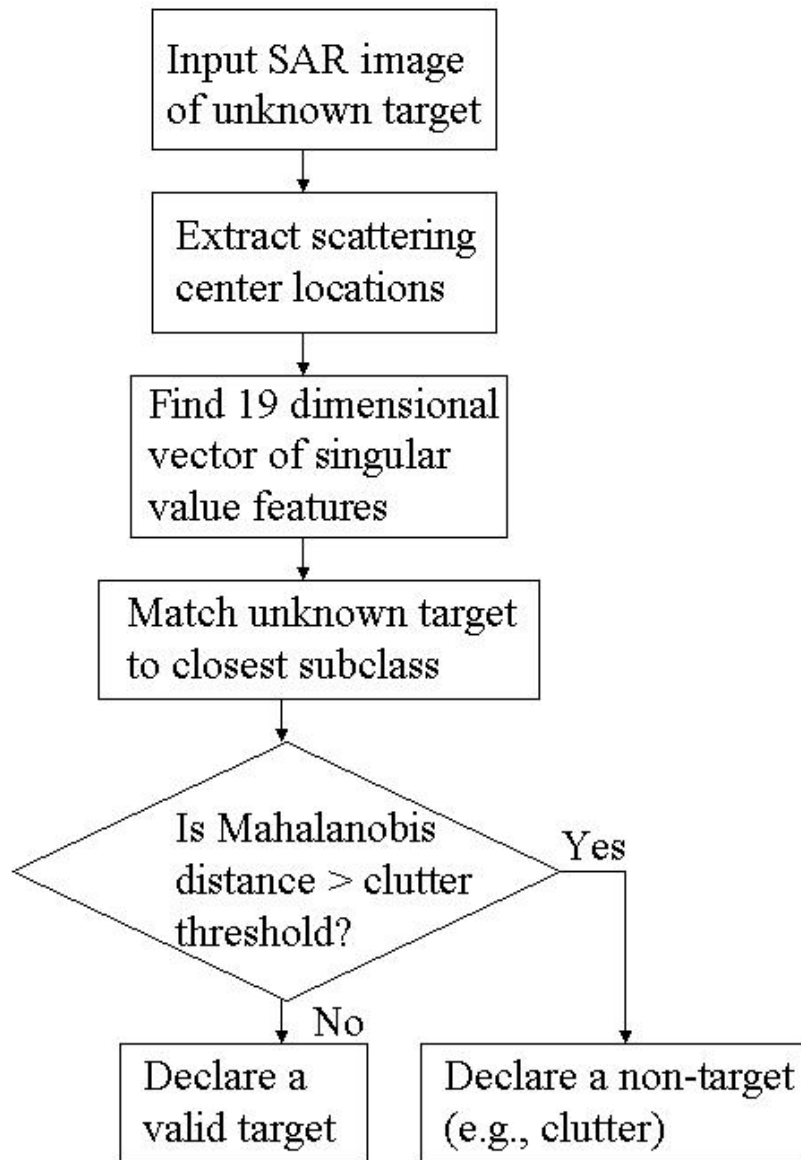


Figure 11. Flowchart of the classifier decision process. First, scattering centers are extracted from the radar image. Next, the Mahalanobis distance of a SVD feature set for an unknown target is found to the closest subclass means in the target database. Finally, if this distance is less than the clutter threshold, a valid target is declared.

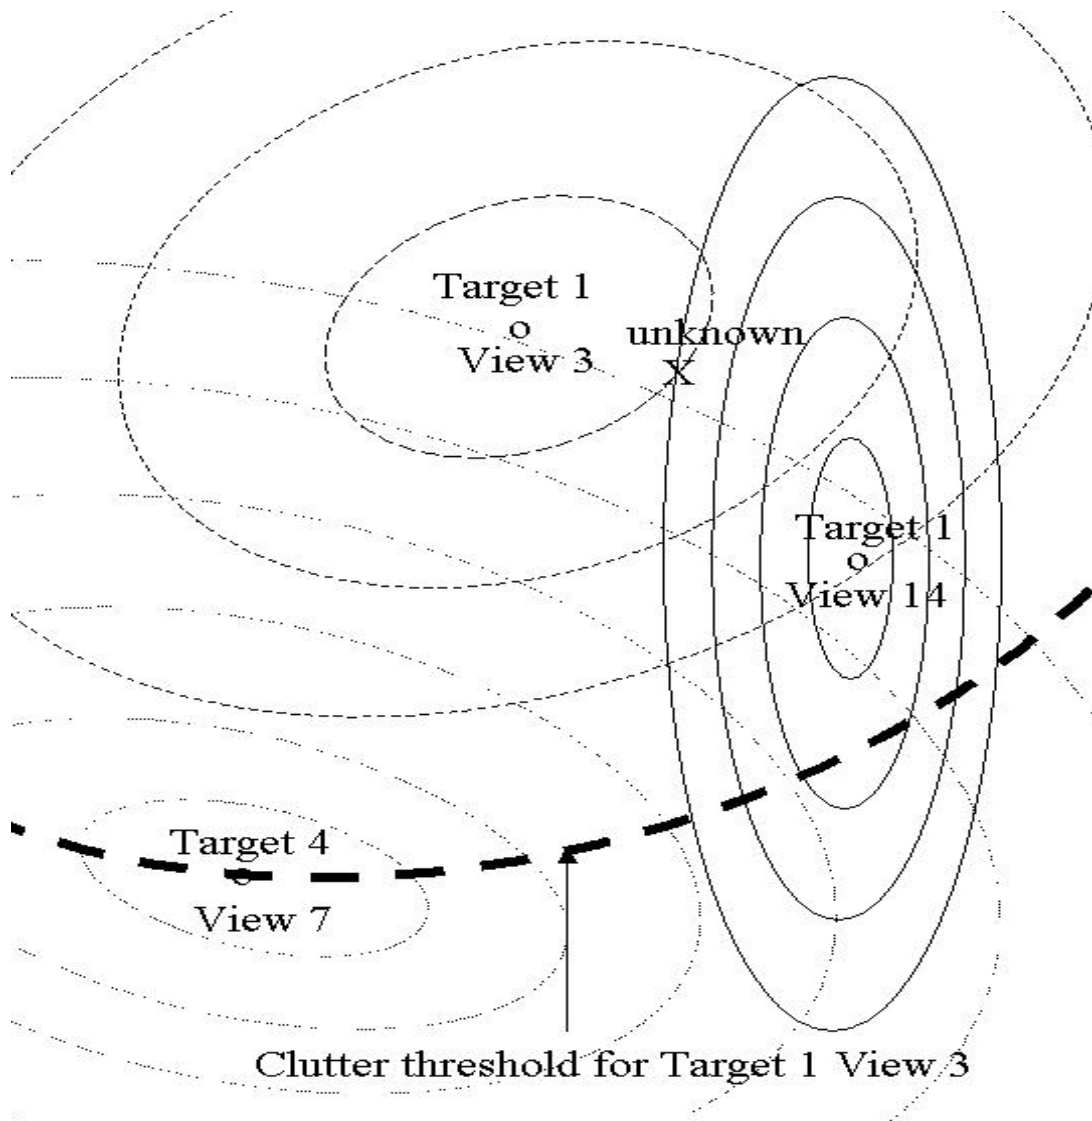


Figure 12. Example of decision and clutter thresholds. An unknown target point (X) and the three closest subclass mean points (o's) are shown, where to enable illustration the points have two dimensions instead of nineteen dimensions. The subclass ellipses indicate lines of equal probability density for the covariance matrix of the indicated view of each target and represent integer multiples of one standard deviation. Any unknown point is matched to the nearest subclass using Mahalanobis distance, i.e., Euclidean distance after the nineteen coordinate axes are rotated and scaled so that the subclass covariance matrix is the unit matrix. If this distance is greater than the Mahalanobis distance to the closest subclass belonging to a different target, target identification is unreliable.

5 Classifier Test

The best test of a classifier is its performance in classifying new (not used in training) images of both targets and “confusers”, where the latter are objects that differ materially (i.e., by more than added noise) from any target. However, only different aspect angles of each target and not different obscurations, clutter background, etc., are available. Nevertheless, there are established methods for replicating “real world” conditions that degrade the original target images.

5.1 Replicating “real world” conditions

To replicate “real world” conditions and exhaustively test the classifier, both white noise and obscuration masks are added to the sets of scattering centers withheld for testing. White noise models the sum of many random variables, and an obscuration mask removes three-dimensional scattering centers from the target model in blocks (which simulates placing the target behind objects such as trees or partially covering the target with non-reflective tarpaulins).

White noise is added to the target with variances up to those of the target coordinates. Obscuration filters are added that mask 0% to 60% of the scattering centers of the target along an arbitrary coordinate axis and at an arbitrary position on the target. Specifically, an equal number of each set of target views is sorted in increasing order for each of the three coordinates of the scattering centers (a third are sorted by x, a third are sorted by y, and a third are sorted by z). At a random position along each realization of a particular target view, a continuous block of scattering centers is removed. For example, if a target view has 1,000 scattering centers and an obscuration level of 20% is desired, 200

consecutive scattering centers are deleted starting with anywhere from the 1st to the 801st scattering center.

5.2 *Calculating confusion matrices*

After degrading the withheld test data, Monte Carlo experiments are performed. The subclasses to which the withheld data is matched using Mahalanobis distance are calculated, and the performance of the classifier is assessed using confusion matrices [18].

A single 288 by 288 matrix describes all test data (degraded by noise or obscuration for input to the Mahalanobis distance classifier), since eight targets are separated into thirty-six subclasses for a total of 288 subclasses. Each element in this confusion matrix has the number of classified test instances for which its row-classified target is its column actual target. For example, if element (5, 8) of the confusion matrix is three, then for all test instances originally in subclass 5 after being subjected to a set of “real world” conditions, three instances are identified as matching subclass 8 the best. The more diagonal the confusion matrix, the better the classifier performance, since each off-diagonal term represents incorrect identification.

A confusion matrix with all 288 subclasses is of limited interest, since only identifying the correct target, and not its aspect angle, is necessary. Thus, the 288 by 288 confusion matrix is reduced to a more tractable eight by eight confusion matrix by summing all 64 of the 36 by 36 submatrices (see Tables 2 through 5). The only information lost in this process is the aspect-angle view of the target under “real world” conditions and the aspect-angle view of the target to which it is matched.

Table 2. Confusion Matrix for 10% Obscuration

	1	2	3	4	5	6	7	8
1	172	0	0	0	8	0	0	0
2	0	179	1	0	0	0	0	0
3	0	3	175	0	0	0	0	2
4	0	0	0	180	0	0	0	0
5	0	0	0	0	180	0	0	0
6	0	0	0	0	0	180	0	0
7	0	0	0	0	0	0	180	0
8	0	0	0	0	0	0	0	180

Table 3. Confusion Matrix for 15% Obscuration

	1	2	3	4	5	6	7	8
1	156	0	0	0	23	0	1	0
2	0	172	7	0	1	0	0	0
3	0	3	170	0	1	0	0	6
4	0	0	0	180	0	0	0	0
5	0	0	0	0	180	0	0	0
6	0	0	0	0	0	180	0	0
7	0	0	2	0	0	0	178	0
8	0	2	0	0	1	0	0	177

Table 4. Confusion Matrix for 25% Obscuration

	1	2	3	4	5	6	7	8
1	97	0	10	0	71	2	0	0
2	1	137	22	0	8	6	1	5
3	0	31	128	0	11	5	0	5
4	0	0	0	180	0	0	0	0
5	0	7	8	0	163	2	0	0
6	0	0	0	0	0	180	0	0
7	2	16	40	0	10	0	87	25
8	0	23	7	0	6	1	6	137

Table 5. Confusion Matrix for 40% Obscuration

	1	2	3	4	5	6	7	8
1	62	1	22	0	47	35	13	0
2	4	85	34	0	13	10	4	30
3	1	43	91	0	8	34	0	3
4	0	0	0	180	0	0	0	0
5	0	59	20	0	55	46	0	0
6	0	0	0	0	0	180	0	0
7	13	21	89	0	21	7	16	13
8	0	70	32	0	5	8	3	62

5.3 Calculating figures of merit

From each confusion matrix a probability of correct identification (PCID) and a probability of false identification (PFID) for each target is calculated.

The PCID for each target is found from the eight rows of the confusion matrix. The sum of the elements in a particular row represents all instances of a single target subject to “real world” conditions, and the column of each row element indicates how many times this target is matched to each of the eight targets. The diagonal element represents a correct identification, so the PCID is defined as the diagonal element divided by the sum of the row elements.

The sum of the elements in a particular column represents all instances in which any target subject to “real world” conditions is identified as a particular target. All off-diagonal elements in the column represent false identification, so the PFID is defined as the sum of all diagonal elements divided by the sum of the column elements.

A particular set of “real world” conditions generates a single confusion matrix, and from this confusion matrix a PCID and PFID are calculated for each of eight targets. A single parameter (noise power or obscuration level) is varied to generate multiple confusion matrices. Converting each confusion matrix into eight PCIDs and eight PFIDs allows the production of graphs of PCID and PFID versus the parameter for each target. Such graphs are displayed in Figures 13 to 16.

5.4 Test results

Adding white noise of variance up to the variance of the test data [Signal to Noise Ratio (SNR) = 1 = 0 dB] to the data withheld for testing does not adversely affect classifier performance because the subclass confusion matrix remains diagonal. Not only are all targets correctly identified. Also, each Mahalanobis distance from test data with added noise of SNR = 0 dB to the correct subclass is less than one ten thousandth of the clutter rejection threshold.

However, obscuring the test data quickly degrades classifier performance as shown in Figures 13 to 16. With one exception, when the test data is obscured by 15%, the PCID for all targets is greater than 95% and the PFID for all targets is less than 5%. With the same exception, when the test data is obscured by 20%, the PCID for all targets is greater than 80% and the PFID for all targets is less than 20%. In summary, when 20% of the test targets are obscured, the classifier still correctly classifies over four out of five cases.

5.5 Comparisons with existing classifiers

The Mahalanobis distance classifier using singular value features shows significant immunity to white noise. For targets subject to obscurations of up to 20%, the classifier performance degrades, but it is still acceptable. An existing two-dimensional classifier has difficulty identifying targets that are only 10% obscured, and it requires over fifty times the memory to store each target and over sixty five times as much computer processor time to make a single comparison as the classifier proposed here [53]. These dramatic improvements can be attributed in part to more stable scattering centers and the data compression enabled by this stability.




Figure 13. Plot of Probability of Correct Identification (PCID) and Probability of False Identification (PFID) as a function of obscuration level for targets 1 & 2. In this figure and the following three figures, the confusion matrices for nine obscuration levels are used to form nine data points on a graph which shows the probability that an obscured version of each target is correctly identified for a particular obscuration level (PCID). Also shown is the probability that an obscured version of another target is misidentified as each of eight targets (PFID).




Figure 14. Plot of PCID and PFID as a function of obscuration level for targets 3 & 4.





Figure 15. Plot of PCID and PFID as a function of obscuration level for targets 5 & 6.





Figure 16. Plot of PCID and PFID as a function of obscuration level for targets 7 & 8.

6 Contributions

This section summarizes contributions of the research reported here, beginning with valuable results not used in the final classifier design and ending with a general purpose "recipe" that characterizes this design.

In the pursuit of a classifier that uses invariant features, valuable information about the relationship between sensor resolution and target identification performance as well as the relationship between Procrustes distance and Kendall shape space was obtained as described in Appendices 1 and 2. This information is not directly related to the singular value feature classifier described and evaluated in Chapters 3 through 5.

Two contributions in the extraction of scattering centers make the design of an invariant three-dimensional classifier practical: the iterative model-building algorithm described in Chapter 2 and the modification of existing commercial software (XPATCH) to extract scattering center models consistent with future capabilities (e.g., 3D MAGI) described in Section 4.3. An efficient model-building algorithm allows a database of targets for the classifier to be constructed with minimal time and cost. The modification of existing commercial software enables rapid testing of the classifier using existing scattering center models.

The primary contribution is a method for classifying targets (including military targets) using invariant features. This method chooses a space that is independent of any translation or rotation of the target and that is also independent of the number of scattering centers that characterize the target. The classifier compresses the input data by

factor of ten million to one, identifies targets correctly in realistic noise environments, and performs acceptably when targets are obscured.

An algorithm that classifies any object which can be represented as a “cloud of points” has applications beyond the field of ATR. Although the algorithm is used here to identify military targets from scattering center models of radar targets, it may be extended to more generic classification problems. The following “recipe” reviews the steps involved in the algorithm.

6.1 Extract “cloud of points” model

Extract the three-dimensional coordinates of all the points necessary to represent each object as series of aspect-angle views. Specifically, let the total number of points needed to represent each aspect-angle view be N and represent each aspect-angle view by an $N \times 3$ matrix, where the columns are the x , y , and z coordinates.

In the radar classification example, three-dimensional scattering centers are extracted every tenth of a degree in aspect angle for each of eight military targets. A modified version of XPATCH is used to extract scattering centers from existing target models, where the scattering centers are constrained to persist for twenty degrees of aspect angle.

6.2 Generate singular value features

Convert each aspect angle view into a vector of nineteen singular value features. Specifically, find the singular values of the original $N \times 3$ matrix after the means are subtracted from each column. Then find the six singular values of the $N \times 6$ matrix with zero-mean columns x^2 , y^2 , z^2 , xy , xz , and yz , and find the ten singular values of the $N \times$

10 matrix with zero-mean columns x^3 , y^3 , z^3 , x^2y , x^2z , xy^2 , xz^2 , yz^2 , y^2z , and xyz . Finally, concatenate all singular values into one length-nineteen feature vector.

6.3 *Characterize aspect-angle subclasses*

Divide the series of aspect angle views for each object into multiple subclasses, where each subclass has as many common points as possible (this commonality is necessary to model the obscuration of points as an object rotates). Withhold 5% of the aspect-angle views, uniformly distributed throughout the subclass, for later testing. Find the mean and covariance matrix of the remaining 95% of the aspect-angle views.

In the radar classification example, aspect-angle views at 1.0, 3.0, 5.0, 7.0, and 9.0 degrees within each 10-degree-wide subclass are withheld, and a mean and covariance matrix are found for the 95 remaining aspect-angle views.

6.4 *Find clutter threshold and classify test points*

Find the Mahalanobis distance from each subclass mean to the closest subclass mean of a different target. Declare this distance to be the clutter threshold, i.e., any unknown is declared a non-target if the Mahalanobis distance to the closest subclass exceeds this threshold.

Classify test points with white zero-mean Gaussian noise added to the coordinates of the $N \times 3$ object matrix with noise variance up to that of the variance of the coordinates themselves (this maximum models a signal to noise ratio of 1).

Classify test points after obscuring the object matrices up to 60% by removing continuous, randomly located and randomly directed blocks of coordinates from the

object matrix. Specifically, order the object points by each of the three coordinates and remove a block of points starting at an arbitrary row in the object matrix.

6.5 *Generate confusion matrices*

Let the row index describe the actual subclass in which the test data belongs and the column index describes the subclass in which the test data is classified.

Generate a square confusion matrix that describes test data classification for each noise level and each obscuration level. Reduce this confusion matrix to a smaller matrix that shows how each target is classified independent of aspect-angle subclass. In the radar classification example, the original 288 by 288 (36 subclasses for each of 8 targets) confusion matrix is reduced further to an 8 by 8 target confusion matrix.

6.6 *Plot PCID and PFID for each target*

Plot PCID and PFID for each target with noise and with obscuration, where PCID is the diagonal element in the column of the target divided by the sum of the corresponding row elements, and PFID is the sum of the off-diagonal column elements divided by the sum of the column elements for the target.

For the radar classification example, the subclass confusion matrix for each target with noise is diagonal, so the probability of correct identification (PCID) is one and probability of false identification (PFID) is zero, and no plots are necessary. However, for each target with obscuration, a plot of PCID and PFID is generated for each obscuration level (see Figures 13 to 16).

7 Conclusions and Recommendations

This dissertation has described methods for building scattering center models of SAR targets, for using these models to train a classifier, and for employing this classifier to identify unknown targets and estimate their aspect angle. Key conclusions and recommendations that follow from these methods are presented in this chapter.

7.1 Conclusions

Target classification using three-dimensional SAR data (which requires complicated relative motion between target and sensor) has improved performance relative to existing two-dimensional SAR ATR algorithms. For example, momentary obscurations of the target or changes in the target (e.g., due to rotation of a turret) typically do not unduly confuse target tracking and identification.

7.1.1 Feature selection

Although Procrustes distance and shape space metrics are promising for ATR, their application to objects with an unequal number of scattering center points is problematic. A feature set that solves this correspondence problem and that is invariant to object translation and rotation consists of the singular values of the object matrix augmented by the singular values of its second and third order monomial expansions. A classifier that uses this singular value feature set displays significant immunity to white noise (added to scattering centers extracted from radar returns) for targets of military interest. Also, when such targets are arbitrarily obscured, the classifier that uses this singular value feature set performs adequately.

7.1.2 Model building

An algorithm that efficiently constructs three-dimensional target models is essential for the development of the classifier, since model construction using airborne sensor measurement of all targets over the entire viewing sphere would be prohibitive in time and cost. The iterative extraction algorithm described here locates individual scattering centers from complex radar returns when the relative position between the target and sensor is known. This extraction algorithm determines all three coordinates of each scattering center.

7.1.3 Classifier design

Separating a target class into multiple subclasses representing narrow windows of aspect angle performs well due to a lack of correlation between views of the same target in different windows. The tests conducted here indicate that each target may be represented by 36 ten-degree subclasses. The correlation of the scattering center points within each window allows each such set of points to be modeled by a normal distribution and thus characterized by a mean vector and a covariance matrix. This result motivates the use of a Mahalanobis distance metric and enables the specification of practical thresholds for avoiding unreliable decisions in the presence of noise and clutter.

7.1.4 Classifier performance

The classifier developed here represents a long three-column object matrix (typically over a thousand rows) as a nineteen-dimensional feature vector. Nevertheless, this classifier functions well for noisy test data and test data obscured by as much as 20%.

7.2 Recommendations

Key suggestions for further development of the methods described and evaluated in this dissertation are presented in this section.

7.2.1 Comprehensive testing

More comprehensive testing of the entire classifier training and testing process is needed. Future research should focus on verification using “more difficult” targets and types of noise and obscuration. A classifier that performs well for a large set of realistic scenarios should be considered for use in a fielded system.

7.2.2 Obscured target improvements

The performance of the classifier proposed here could be improved for obscured targets. Each target could be segmented into several sections, and the singular value features of the sections could be stored as a function of aspect angle. If any obscured target does not match any unobscured target adequately, its features could be compared in Mahalanobis distance to features of the sections to identify a better match. This segmentation technique should contribute to the solution of challenging obscuration problems.

7.2.3 Improved classification algorithm with confidence measures

Although matching an unknown target point to the closest subclass mean in Mahalanobis distance adequately classifies the eight military targets considered here, further improvements that specify a posterior probability for each target would lead to associated confidence measures.

The Mahalanobis distance to each target subclass mean is calculated by the current algorithm. The posterior probability that an unknown point belongs to subclass i is $\exp(-r_i^2/2)/\sum_j \exp(-r_j^2/2)$, where r is the Mahalanobis distance to the mean of subclass i and the summation is over all subclasses. The posterior probability that an unknown point belongs to a specified target is the sum of all the posterior probabilities that an unknown point belongs to each subclass associated with this target. These specifications yield, as required, unit total posterior probability for all targets (or, alternatively, for all subclasses).

A standard confidence measure is the posterior probability associated with each target. The decision to match an unknown point to a particular target improves as this confidence measure approaches one. A practical implementation would rank the posterior probabilities from largest to smallest for all targets and select the target associated with the largest posterior probability. If the next largest posterior probability is close to the largest posterior probability, the classifier could declare uncertainty and indicate the alternative targets as well.

Appendix 1

Resolution Study

This appendix describes a resolution study [57] that uses eight 1700x1700 pixel images of four nearly identical targets, each at two slightly different aspect angles (Figure A1). The raw SAR data is converted to magnitude-only form after all processing is completed.

The four targets are standard mid-sized cars: Grand Prix, Grand Marquis, Intrepid, and Lumina. In each scene all targets (cars) have the same physical orientation, so no rotation is needed for target comparison within a scene. All eight scenes are within 30 degrees in azimuth, but the illumination of the targets in each scene is unique. Also, all targets are spaced far enough apart to ensure that the images and their shadows are distinct.

1 Methodology

1.1 Registration

A 256x256 chip of the first target and its shadow is chosen as a reference. This reference chip is shifted across the other three target chips in the same scene and the lowest Mean Squared Error (MSE) match for each target is recorded. Each chip is normalized to unit energy to minimize differences due to angle of illumination and clutter. The chips in the reference image are then registered to the chips of the same vehicle in the seven remaining images. The registered images are normalized to unit energy and stored in large four-dimensional arrays.

1.2 Decreasing resolution by “coarsening”

Standard filtering techniques and sub-sampling artificially decrease the resolution of all the images to four new resolutions which are one half, one fourth, one eighth, and one sixteenth times the resolution of the original images in each dimension. These powers of two make the sub-sampling simple: to sub-sample by a factor of two, retain all the odd pixels in each dimension. The term “coarsening” refers to this filtering and sub-sampling process.

Starting with the original complex 256x256 images, coarsening by a factor of N involves first choosing a twenty tap equiripple filter that passes the lower $1/N$ of the complex frequency spectrum. This equiripple filter minimizes the ringing and aliasing that would otherwise occur by truncating in the frequency domain to implement ideal low-pass filters. The Parks-McClellan transform takes this one-dimensional filter and transforms it into an appropriate cone-shaped two-dimensional filter (see Figure A2). This two-dimensional filter is convolved with the complex-valued original high-resolution image with the stipulation that the output is the same size as the input.

After filtering, the resulting image is sub-sampled by a factor of N in both dimensions. The set of chips now includes 256x256, 128x128, 64x64, 32x32, and 16x16 pixel images. Sub-sampling is critical: a low-resolution sensor will not only lose the high frequency information of the high-resolution sensor, but it will also produce fewer pixels.

1.3 Confusion matrix

A 32x32 confusion matrix is generated for each resolution (each car illuminated at a different azimuth angle is treated as a separate target). The illumination profile of the same target changes drastically over a small angle when illuminated by a radar system.

A MSE is calculated for each target compared to the database of all thirty-two target images at each resolution. Note, however, that the phase is important to the filtering and subsampling process: had phase been discarded prior to filtering, valuable information would have been lost, and because comparisons are made from one resolution to another, the MSEs are, by definition, normalized by the number of pixels in the chip. All images are normalized to unit energy, and this normalization takes into account the number of pixels in the image.

1.4 Spatial domain interpolation

Lower resolution data should, on average, have a higher MSE when comparing a target to the thirty-one other targets. Using interpolation in the spatial domain, artificial higher-resolution data is created. Theoretically, once higher frequency information is eliminated it cannot be recovered, but interpolation creates an estimate of the lost data. This artificial higher resolution data can be treated as a separate resolution, and thus another confusion matrix is generated for these artificial resolutions.

2 Data

2.1 Average between-class MSE

Each confusion matrix has all zeros on the main diagonal. All off-diagonal elements correspond to a comparison of a target image to a “confuser”. The average of the off-diagonal elements is chosen as a simple metric to compare different resolutions. The larger the average of this off-diagonal element, referred to here as the average between-class MSE, the easier it is to separate these similar targets.

2.2 Probability of false alarm versus MSE threshold

The data is a distribution of MSEs, so a plot of probability of false alarm versus MSE threshold is more descriptive than average between-class MSE (a single number). A MSE threshold is chosen for comparison with all possible “confusers”: any target with an MSE greater than this threshold is correctly classified. As there is only sample of each target, it is not possible to generate Receiver Operating Characteristic (ROC) curves. However, it is possible to express the probability of false alarm (confusing another target with the actual target) as a function of the MSE threshold for all resolutions, including resolutions artificially created using interpolation.

3 Results

As resolution is increased by a factor of two, the average between-class MSE increases by five to ten percent. In this thirty-two target classification problem, the average between-class MSE for each of the resolution cells are as follows: .9186 for one, .8689 for one half, .7747 for one quarter, .6835 for one eighth, and .6340 for one sixteenth times the original resolution. These values are plotted versus the log of the resolution (see Figure A3).

The one quarter resolution is linearly interpolated to yield artificial one half resolution data, which was again interpolated to full resolution data to compare with the uninterpolated one half and one quarter resolution data. The average between-class MSE for interpolated full resolution and one half resolution data is .6011 and .6355 respectively.

The curves of probability of detection versus MSE threshold (see Figure A4) show the differences from one resolution to another. Artificial resolutions created by interpolation are compared to high-resolution and low-resolution images (see Figure A5).

4 *Conclusions*

“Coarsened” images allow five to ten percent less discrimination than the original higher resolution images in this problem. The larger average between-class MSE of the original images facilitates the correct identification of “confusers”. Interpolating the “coarsened” images to an artificially higher resolution degrades target identification performance significantly. Thus for this many-class target identification problem, interpolation is a poor method for reconstructing lost frequency information.

“Coarsening” high-resolution images must remove more image information than clutter and noise at high frequencies, hence the five to ten percent degradation for each halving of the resolution. The “coarsened” images, in addition to being harder to separate using a MSE criterion, are twice as small in each dimension. This compression may justify the decreased discriminating power in some applications.

5 *Recommendations*

Obtaining additional experimental data in the form of multiple samples of the same target at the same illumination angle would allow comparisons of images of the same target type in addition to comparisons between two different target types. ROC curves using a moving MSE threshold could then be generated to provide a more powerful means to compare sets of target images of different resolutions (including artificially interpolated resolutions).

If additional experimental data is unavailable, adding noise to the set of target images with second and possibly third order statistics equivalent to the clutter background of the original images would allow the same powerful means of comparing target identification performance between different levels of resolution, but would depart significantly from a “real world” scenario.



Figure A1. High resolution SAR image of four similar cars: Grand Prix, Grand Marquis, Intrepid, and Lumina. The image is complex-valued, and only its magnitude is presented. Note that the similarity of the four targets makes this a challenging many-class problem.




Figure A2. Frequency domain representation of the two-dimensional equiripple filter with the Parks-McClellan transform that “coarsens” by a factor of two. After this filter is implemented on the complex data with the constraint that the output image be the same size as the input image, the resulting image is subsampled by a factor of two to complete the “coarsening” process.




Figure A3. Plot of the average between-class MSE versus the log of resolution. Every doubling of the resolution of the original high resolution images causes a five to ten percent decrease in the average between-class MSE. This decrease in between-class MSE makes it difficult to separate the true target image from the other 31 target images.




Figure A4. Plot of probability of false alarm versus the MSE threshold for five different resolutions. These resolutions are (from left to right on the graph) one sixteenth, one eighth, one fourth, one half, and one times the original resolution. The probability of false alarm is the fractional chance of confusing another target with the actual target. Each doubling of the resolution shifts the probability of false alarm versus MSE threshold curve farther to the right and decreases the sensitivity of target identification performance to the MSE threshold.




Figure A5. A plot of probability of detection versus MSE threshold for various resolutions. The curves on the graph from left to right are one fourth resolution interpolated to one times the original resolution, one fourth resolution interpolated to one half times the original resolution, one fourth times the original resolution, one half times the original resolution, and one times the original resolution. Interpolating the coarsened image to estimate lost information yields worse target identification performance than the coarsened images themselves.

Appendix 2

Procrustes Distance and Shape Space

Archeologists and biologists devised Procrustes distance as a “one-size-fits-all” metric, and it is a proven metric for classifying objects by shape [11]. Calculating the Procrustes distance between two objects requires representing both objects as sets of points. To compare continuous objects, unique points (landmarks) common to both objects (e.g., the tip of the nose on an animal skull) are needed.

1 Relevance to ATR problem

In the three-dimensional ATR problem, an object is represented as a finite collection of points. Radar ray-tracing algorithms also represent an object as a finite collection of point scattering centers. The shape of an object is defined in terms of properties that are independent of any translation, rotation, or scaling of the object coordinates.

If one knows the a priori probabilities of the targets in Euclidean (object) space and desires to design a classifier in a shape space, it is necessary to know the corresponding a priori probabilities in shape space. If one accepts a uniform distribution for possible targets in object space (i.e., all possible shapes are equally likely), then decision boundaries in shape space may be calculated.

2 Shape space experiments

Two useful experiments are performed here. First, assuming a uniform distribution of targets in object space, the resulting distribution in shape space is characterized. Second, for a simple example, i.e., triangles in two dimensions, an isometric relationship between Procrustes distance and shape space is shown. Although it is difficult to characterize the distribution of object coordinates in nature, a uniform distribution is a good starting point.

The triangles considered here are the simplest of objects, but should provide insight into how to handle more complicated objects.

2.1 Characterizing shape space

One thousand triangles are randomly generated such that each corner of the triangle is located inside the box bounded by $(0, 0)$, $(0, 1)$, $(1, 1)$, and $(1, 0)$. Figure A6 shows the first ten randomly generated triangles. The dimension of shape space for three points in two dimensions is two: initially there are six coordinates, but included are two dimensions for location, one dimension for uniform scale, and one dimension for rotation. Therefore, two degrees of freedom remain [11]. Typically two points of the triangle (known) are mapped to two specific (known) points in the shape space allowing the recreation of the two by two permutation matrix. Now the third point of the triangle multiplies the permutation matrix, and the coordinates of this transformed point are the shape space coordinates.




Figure A6. Random triangles in object space. Here selecting three points randomly in the box bounded by $(0, 0)$, $(0, 1)$, $(1, 1)$, and $(1, 0)$ forms each of ten triangles. To classify these triangles by shape, it is necessary to find their distribution in two-dimensional shape space.

Although there are many possible definitions of shape space, Kendall coordinates are a popular choice. In these coordinates the first point is mapped to the point $(-1/2, 0)$, the second point is mapped to the point $(1/2, 0)$, and the coordinate of the third point is

$$(u_3^K, v_3^K) = \frac{2}{\sqrt{3}} \frac{z_3^O - \frac{1}{2}(z_1^O + z_2^O)}{z_2^O - z_1^O} \quad (\text{A-1})$$

where superscripts K designate Kendall coordinates, superscripts O designate object coordinates, and the coordinates are treated as complex numbers [11]. With this particular transformation to Kendall shape space, the shape of any triangle is represented as a single two-dimensional vector or complex number. Objects with more than three points require mapping each point after the first two to a point in two dimensions.

If triangles are uniformly distributed as shown in Figure A6, to show the distribution of the shape space coordinates (the coordinates of the third point of the triangles calculated from Equation (A-2)) are bivariate, Gaussian distributions (see Figure A7) requires hypothesis testing. To test the hypothesis that the distribution $F(x)$ of a random variable x equals a given function $F_0(x)$, the hypothesis that $F(x) = F_0(x)$ at a set of $m - 1$ points $\{a_1, a_2, \dots, a_{m-1}\}$ is examined as shown in Equation (A-2) using a Chi-Square test with $m - 1$ degrees of freedom.

$$\begin{aligned} H_0 : F(a_i) &= F_0(a_i), 1 \leq i \leq m-1 && \text{Null hypothesis} \\ H_1 : F(a_i) &\neq F_0(a_i), \text{some } i && \text{Alternate hypothesis} \end{aligned} \quad (\text{A-2})$$

The procedure involves calculating the difference between the number of samples in an interval, squaring the result, and dividing by the expected result in the interval. The resulting values are summed over all the intervals to obtain a number for use with a Chi-

Squared table. If the weighted sum squared difference is less than the number from the table, the null hypothesis is accepted [35].

2.2 Relating shape space to Procrustes distance

A triangle with Kendall shape coordinates (0, 1) is compared to all possible triangles in the shape space. The Procrustes distance between two triangles, represented as complex vectors w and y , respectively, is calculated as a linear regression of y on w

$$d_F(w, y) = \left\{ 1 - \frac{y^* w w^* y}{w^* w y^* y} \right\}^{1/2} \quad (\text{A-3})$$

where F denotes the full Procrustes metric, the $*$ corresponds to the conjugate transpose, and the original objects have zero mean [11]. The one thousand triangles are projected into Kendall shape space, and their Kendall coordinates are shown as a scatter plot in Figure A7.

3 Data

For better representation of the probability density function, the shape space is segmented into uniform squares and converted to a discrete probability density function through a histogram. The best normal distribution fitting this data has a mean of (0.0961, 0.0191) and a standard deviation of (1.9745, 1.9375). To perform hypothesis testing, another histogram is generated for comparison using the postulated bivariate normal probability density function.

The shape space may be uniformly sampled to generate a mesh plot of Procrustes distance versus the shape space coordinates of an arbitrary triangle compared to a reference triangle. Figures A8 and A9 show that if reflections are ignored and if only the

half-plane containing the reference triangle is considered, Procrustes distance is isomorphic to the Kendall coordinate due to the lack of local minima.

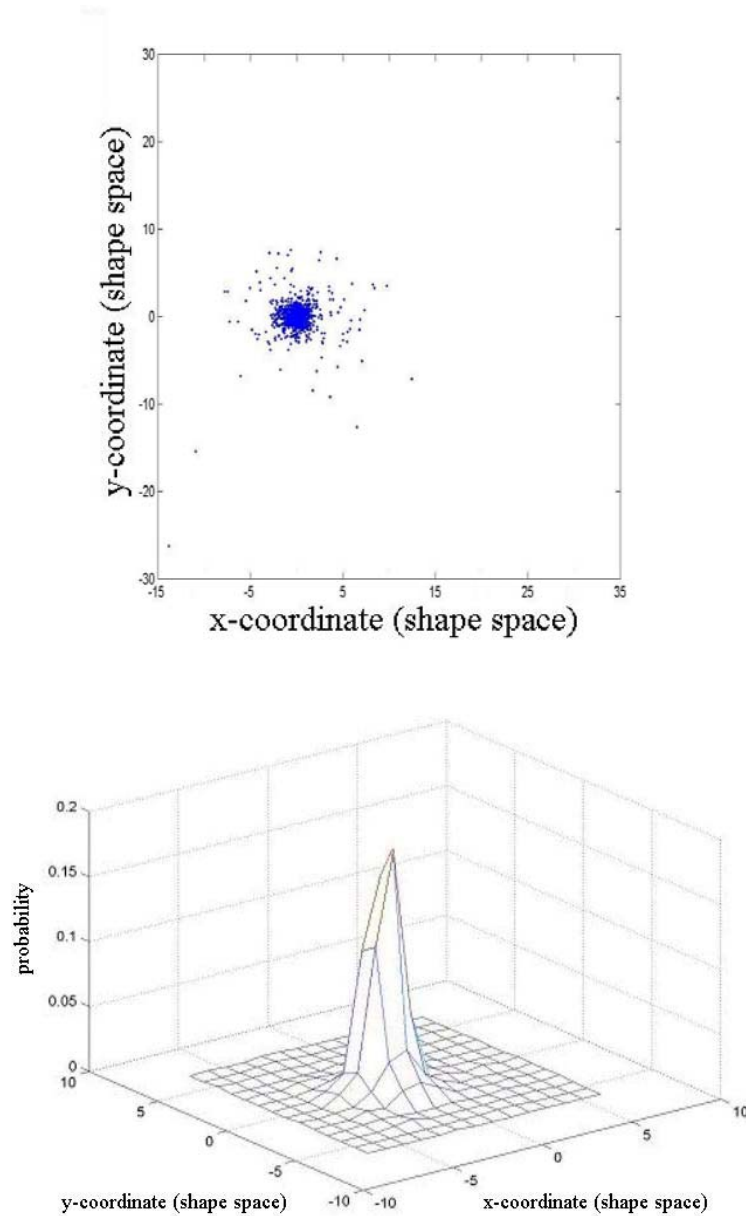


Figure A7. Scatter plot and approximate probability density function of a thousand triangles in shape space. Here triangles that are uniformly distributed in Euclidean space are mapped to Kendall shape space, where for each triangle one point is mapped to the point $(-1/2, 0)$, a second point is mapped to the point $(1/2, 0)$, and the coordinate of the third point is found from Equation (A-1). The data suggest an independent (no correlation between horizontal and vertical axis) bivariate normal probability density with mean $(0, 0)$ and with a standard deviation for both variables of approximately two. Standard hypothesis testing shows that this distribution is Gaussian.




Figure A8. Plot of Procrustes distance in Kendall shape space. Procrustes distance from Equation (A-3) is found for a reference triangle whose Kendall coordinate is (0, 1). There is an obvious minimum at (0, 1) (an exact match). Although the Procrustes distance metric equates two objects that differ by any permutation matrix that combines any combination of translation, rotation, or scale, it does not account for reflections. There is a maximum at (0, -0.75) which corresponds to a reflected version of the original triangle. (A better distance metric may be the minimum of two Procrustes distances: the Procrustes distance between the two objects, and the Procrustes distance between one object and the reflection of the other object.)




Figure A9. Expanded plot of Procrustes distance in shape space. Here the lack of correlation between Procrustes distance and distance in shape space is shown. As the reference triangle is compared to “stretched” triangles (y increases from 1), Procrustes distance is least sensitive to change in shape space. As the reference triangle is “warped” (x increases or decreases from 0), Procrustes distance is more sensitive to change in shape space. Finally, as the reference triangle is “shrunk” (y decreases from 1), Procrustes distance is most sensitive to change in shape space.

4 Results

The probability density function for uniformly distributed triangles mapped to the shape space is approximately a diagonal-covariance bivariate normal distribution, i.e.,

$$f(x)f(y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left\{-\frac{1}{2}\left[\left(\frac{x-\mu_x}{\sigma_x}\right)^2 + \left(\frac{y-\mu_y}{\sigma_y}\right)^2\right]\right\} \quad (\text{A-4})$$

where $(\mu_x, \mu_y) \cong (0, 0)$; $\sigma_x \cong 2$; $\sigma_y \cong 2$ [35].

The Chi-Square test to accept x and y as normally distributed with mean 0, standard deviation 2, and a diagonal covariance matrix requires accepting three null hypotheses that are tested using the sum squared difference q: (1) the x coordinates have a Gaussian distribution with mean 0 and standard deviation of 2, (2) the y coordinates have a Gaussian distribution with mean 0 and standard deviation of 2, and (3) the x and y coordinates are independent

$$q = \sum_{i=1}^{m-1} \frac{\{k_i - [F_O(x_{i+1}) - F_O(x_i)]\}^2}{F_O(x_{i+1}) - F_O(x_i)} \quad (\text{A-5})$$

where k_i is actual data and $F_O(x)$ is the postulated CDF (note $F_O(x_{i+1}) - F_O(x_i)$ are the expected values of the pdf integrated over the interval x_i to x_{i+1}). Rather than show that the x and y coordinates are independent directly, the hypothesis that the angle $\arctan(y/x)$ is uniformly distributed between 0 and π is tested instead.

With 95% confidence, the test for a normal distribution in the x direction yields a q of $6.98 < 38.89$, the test for a normal distribution in the y direction yields a q of $2.75 < 38.89$, and finally the test for a uniform $\arctan(y/x)$ yields a q of $4.55 < 16.9$, where the larger quantity in each case is the Chi-Squared statistic. These results validate all three null hypotheses and verify that the distribution of the Kendall coordinates is described by Equation (A-4) (with an acceptable alpha error of 5%).

To verify that the Kendall coordinates and Procrustes distance are isometric (ignoring reflections), all points within the half-plane are tested to insure that there are no local minima. Using the complex expression for Procrustes distance in Equation (A-3), there is no point in the area of interest where f_x and f_y are both zero and the Hessian $(f_{xx} * f_{yy} - f_{xy}^2)$ is positive, where the single and double subscripts designate the first and second

partial derivatives, respectively, of the Procrustes distance d_f . Also, all points at which f_x are f_y are zero are such that the only minima is at (0, 1), which specifies the coordinate of the standard object against which the other objects are compared.

5 Conclusions

A uniform distribution of triangles in object space corresponds to an independent bivariate normal distribution in Kendall shape space with mean (0, 0) and the same standard deviation in both coordinate directions. This result is due to the interactions of multiple independent variables in object space that generate the shape space coordinate in Equation (A-1), which indicate that the Central Limit theorem applies so that the Kendall shape coordinate is approximately normal [35]. For a large variety of targets in the object space the Central Limit theorem should hold and the shape space distribution should be close to normal. For example, if triangular facets from large objects are used, the variety of facets for all targets should be sufficient to generate a normally distributed shape space.

Procrustes distance is, by definition, independent of rotation, translation, and scale, and objects in the shape space are equivalent over these transformations. As expected, distance in shape space and Procrustes distance are isometric if reflection is ignored. However, there is no simple Euclidean distance relationship between Kendall coordinates and Procrustes distance.

6 Correspondence problems

A set of scattering centers is an efficient way to represent any complex radar return. If a set of scattering centers for a target is projected to shape space, all translations, rotations, and scalings of the target are equivalent. If Procrustes distance is chosen as a

metric for comparing unknown to known targets, then testing all possible combinations of rotation, translation, and scale of the unknown targets can be avoided.

However, in general two targets compared in an ATR algorithm have a different number of points with no known correspondence. This lack of correspondence may lead to an intractable combinatorial problem; e.g., there are about 100 trillion possible ways to match a twenty-point object to a thirty-point object. An approach in which order is unimportant can reduce this number to about 30 million possible ways, but this result is still not satisfactory. Unless there is a way to pare down or efficiently search the set of all possible twenty-point matches in the last example, a metric that preserves the desired translation, rotation, and scale equivalence and that avoids the correspondence problems is required.

Bibliography

- [1] Arnold, G., R. Claypoole, Jr., V. Velten, and K. Sturtz, "Synthesizing Invariant 3-D Rigid Scattering Centers," *Proc. SPIE*, Vol. 4382, pp. 395-403, Orlando, FL, 2001.
- [2] Arnold, G. and K. Sturtz, "Complexity Analysis of ATR Algorithms Based on Invariants," *Proc. IEEE Computer Vision Beyond the Visible Spectrum: Methods and Applications Workshop*, pp. 27-36, Hilton Head, SC, Jun 2000.
- [3] Arnold, G. and K. Sturtz, "Geometric Object/Image Relations for Radar," *Proc. SPIE*, Vol. 4053, pp. 538-545, Orlando, FL, Apr 2000.
- [4] Arnold, G., K. Sturtz, and V. Velten, "Lie Group Analysis in Object Recognition," *Proc. DARPA Image Understanding Workshop*, Vol. 2, pp. 1173-1178, May 1997.
- [5] Arnold, G., K. Sturtz, and V. Velten, "Quasi-invariants of the thermophysical model," *Proc. DARPA Image Understanding Workshop*, pp. 883-892, Monterey, CA, Nov 1998.
- [6] Arnold, G., K. Sturtz, and V. Velten, "Invariants of the LWIR Thermophysical Model," *Proc. IEEE Computer Vision Beyond the Visible Spectrum: Methods and Applications Workshop*, pp. 49-58, Fort Collins, CO, Jun 1999.
- [7] Arnold, G., K. Sturtz, and V. Velten, "Quasi-Invariants of the Thermophysical Model," *Proc. DARPA Image Understanding Workshop*, Vol. 2, pp. 883-894, Monterey, CA, Nov 1998.
- [8] Arnold, G., K. Sturtz, V. Velten and N. Nandhakumar, "Dominant Subspace Invariants," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 22, pp. 649-662, Jul 2000.
- [9] Binford, T. and T. Levitt, "Quasi-invariants: Theory and exploitation," *Proc. DARPA Image Understanding Workshop*, pp. 819-830, Washington DC, 1993.
- [10] Clark, L. and V. Velten, "Image Characterization for Automatic Target Recognition Algorithm Evaluations," *Optical Engineering*, Vol. 30, pp. 147-153, Feb 1991.
- [11] Dryden, I. and K. Martia, Statistical Shape Analysis, John Wiley and Sons, 1999.
- [12] Gauder, M., V. Velten, L. Westerkamp, J. Mundy, and D. Forsyth, "Thermal Invariants for Infrared Target Recognition," *Proc. Third Automatic Target Recognizer Systems and Technology Conference*, Jun 1993.

- [13] Gerry, M., L. Potter, I. Gupta, and A. Van Der Merwe, "A Parametric Model for Synthetic Aperture Radar Measurement," *IEEE Transactions on Antennas and Propagation*, Vol. 47, pp. 1180-1181, Jul 1999.
- [14] Healey, G. and D. Slater, "Using illumination invariant color histogram descriptors for recognition," *Proc. IEEE CVPR*, pp. 355-360, Seattle, WA, Jun 1994.
- [15] Healey, G., "Using color to segment images of 3-d scenes," *Proc. SPIE*, Vol. 1468, pp. 814-825, Orlando, FL, 1988.
- [16] Henle, M., Modern Geometries, The Analytic Approach, Prentice-Hall, 1997.
- [17] Meyer, G., S. Gustafson, G. Arnold, "Extracting Models from radar Data for 3-D Target ID," *Proc. SPIE*, Vol. 4727, pp. 38-45, Orlando, FL, Apr 2002.
- [18] Meyer, G., S. Gustafson, G. Arnold, "The Effect of SAR Resolution on Target Identification," *Proc. SPIE*, Vol. 4382, pp. 338-345, Orlando, FL, Apr 2001.
- [19] Meyer, G., S. Gustafson, G. Arnold, "Three-Dimensional Identification of Moving Targets Using Synthetic Aperture Radar Returns", *Proc. JEWG*, San Antonio, TX, May 2002.
- [20] Meyer, G., S. Gustafson, G. Arnold, "Using Procrustes distance and shape space for Automatic Target Recognition," *Proc. SPIE*, Vol. 4667, pp. 66-73, San Jose, CA, Jan 2002.
- [21] Michel, J., N. Nandhakumar, and V. Velten, "Thermophysical Algebraic Invariants from Infrared Imagery for Object Recognition," *IEEE Transactions on PAMI*, Vol. 19, pp. 41-51, Jan 1997.
- [22] Michel, J., N. Nandhakumar, and V. Velten, "Thermophysical algebraic invariants from infrared imagery for object recognition," *IEEE Transactions on PAMI*, Vol. 19, pp. 41-51, Jan. 1997.
- [23] Mundy, J. and A. Zisserman, Eds., Geometric Invariance in Computer Vision, MIT Press, 1992.
- [24] Mundy, J., A. Zisserman, and D. Forsyth, Eds., Applications of Invariance in Computer Vision, *Proc. Second Joint European-US workshop*, Springer-Verlag, 1994.
- [25] Nandhakumar, N., G. Arnold, J. Michel, G. Tsihrantzis, and V. Velten, "Robust Thermophysics-based Interpretation of Radiometrically Uncalibrated IR Images for

ATR and Site Change Detection," *IEEE Transactions on Image Processing special issue on ATR*, Vol. 6, pp. 65-78, Jan 1997.

- [26] Nandhakumar, N., G. Arnold, J. Michel, G. Tsihrintzis, and V. Velten, "Site Change Detection in IR Imagery from Alpha-Stable modeled Thermophysical Invariant Features," *Proc. SPIE*, Vol. 2645, pp. 122-133, Washington, DC, Oct 1995.
- [27] Nandhakumar, N., V. Velten, and J. Michel, "Thermophysical Affine Invariants for IR Imagery for Object Recognition," *Proc. IEEE Workshop on Physics-Based Modeling in Computer Vision*, pp. 48-54, Boston, MA, Jun 1995.
- [28] Nayar, S., and R.D. Bolle, "Reflectance ratio: A photometric invariant for object recognition," *Proc. ICCV*, pp. 280-284, 1993.
- [29] Nichols, T., J. Thomas, W. Kober and V. Velten, "Interference-Invariant Target Detection in Hyperspectral Images," *Proc. SPIE*, Vol. 3372, pp.176-187, Apr 1998.
- [30] Nichols, T., J. Thomas, W. Kober, D. Arnold and V. Velten, "Canonical Correlation Analysis of LWIR Imagery in the Frequency Domain," *Proc. SPIE*, Vol. 3371, pp. 460-470, Apr 1998.
- [31] O'Neill, B., Elementary Differential Geometry, Academic Press, 1966.
- [32] Olver, P., Applications of Lie Groups to Differential Equations, pp. 73-75, Springer-Verlag, 1986.
- [33] Olver, P., Classical Invariant Theory, London Mathematical Society Student Texts, Vol. 44, Cambridge University Press, 1999.
- [34] Olver, P., Equivalence, Invariants, and Symmetry, Cambridge University Press, 1995.
- [35] Papoulis, A., Probability, Random Variables, and Stochastic Processes, Mc-Graw Hill, 1991.
- [36] Quan, L., "Invariants of six points and projective reconstruction from three uncalibrated images," *IEEE Transactions on PAMI*, Vol. 17, pp. 34-46, Jan 1995.
- [37] Rivlin, E. and I. Weiss, "Local invariants for recognition," *IEEE Transactions on PAMI*, Vol. 17, pp. 226-238, Mar 1995.

- [38] Rothwell, C., D. Forsyth, A. Zisserman, and J. Mundy, "Extracting projective structure from single perspective views of 3d point sets," *Proc. ICCV*, pp. 573-582, Berlin, 1993.
- [39] Ruck, G., Radar Cross Section Handbook, p. 591, Plenum Press, 1970.
- [40] Sattinger, D. and O. Weaver, Lie Groups and Algebras with Applications to Physics, Geometry, and Mechanics, Springer-Verlag, 1986.
- [41] Shashua, A., "Algebraic functions for recognition," *IEEE Transactions on PAMI*, Vol. 17, pp. 779-789, Aug 1995.
- [42] Shashua, A. and M. Werman, "Trilinearity of three perspective views and its associated tensor," *Proc. ICCV*, pp. 920-925, Cambridge, MA, Jun 1995.
- [43] Slater, D. and G. Healey, "Exploiting an atmospheric model for automated invariant material classification in hyperspectral imagery," *Proc. SPIE*, Vol. 3372, pp. 60-71, Orlando, FL, Apr 1998.
- [44] Stiller, P., "General approaches to recognizing geometric configurations from a single view," *Proc. SPIE*, Vol. 3168, pp. 262-273, San Diego, CA, Jul 1997.
- [45] Stiller, P., "Object recognition via configuration of lines," *Proc. SPIE*, Vol. 3454, pp. 76-86, San Diego, CA, Aug 1998.
- [46] Stiller, P., C. Asmuth, and C. Wan, "Invariants, indexing, and single view recognition," *ARPA Image Understanding Workshop*, pp. 1432-1428, Monterey, CA, Nov1994.
- [47] Stiller, P., C. Asmuth, and C. Wan, "Single view recognition – the perspective case," *Proc. SPIE*, Vol. 2826, pp. 226-235, Denver, CO, Aug 1996.
- [48] Stuff, M., "Three-dimensional analysis of moving target radar signals: Methods and implications for ATR and feature aided tracking," *Proc. SPIE*, Vol. 3721, pp. 485-496, Orlando, FL, Apr 1999.
- [49] Stuff, M., PhD dissertation, University of Michigan, to appear, 2002.
- [50] Velten, V., "SAR Image Invariants from 3-D Scattering Centers," *Proc. SPIE*, Vol. 4382, pp. 367-378, Orlando, FL, 2001.
- [51] Velten, V., "Geometric Invariance for Synthetic Aperture Radar (SAR) Sensors: Experimental Results," *Proc. SPIE*, Vol. 3721, pp. 520-531, Orlando, FL, Apr 1999.

- [52] Velten, V., "Geometric Invariance for Synthetic Aperture Radar (SAR) Sensors," *Proc. SPIE*, Vol. 3370, p. 176-187, Orlando, FL, Apr 1998.
- [53] Velten, V., private communication, 2002.
- [54] Velten, V., T. Ross, J. Mossing, S. Worrell, and M. Bryant, "Standard SAR ATR Evaluation Experiments using the MSTAR Public Release Data Set," *Proc. SPIE*, Vol. 3370, p. 566-573, Orlando, FL, Apr 1998.
- [55] Weinshall, D., "Direct computation of qualitative 3d shape and motion invariants," *IEEE Transactions on PAMI*, Vol. 13, pp. 1236-1240, Dec 1991.
- [56] Weinshall, D., Model-based invariants for 3d vision," in *Proc. IEEE CVPR*, pp. 695-696, Cambridge, MA, Jun 1993.
- [57] Weiss, I., "Model-based recognition of 3d objects from one view," *Proc. DARPA Image Understanding Workshop*, pp. 641-652, Monterey, CA, Nov 1998.
- [58] Weiss, I., "Noise-resistant invariants of curves," *IEEE Transactions on PAMI*, Vol. 15, pp. 943-948, Jul 1993.
- [59] Zerroug, M. and R. Nevatia, "Quasi-invariant properties and 3d shape recovery of non-straight, non-constant generalized cylinders," *Proc. IEEE CVPR*, pp. 96-103, Cambridge, MA, 1993.
- [60] Zollars, R. and Z. Hashim, "XPATCH Overview," <http://www.saic.com/products/software/xpatch>, 2002.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 074-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY) 11-04-2003		2. REPORT TYPE Doctoral Dissertation		3. DATES COVERED (From – To) Sep 1999 – Apr 2003		
4. TITLE AND SUBTITLE CLASSIFICATION OF RADAR TARGETS USING INVARIANTS				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Meyer, Gregory, Capt, USAF				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 640 WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/DS/ENG/03-04		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Automatic Target Recognition Branch, Sensor's Directorate, AF Research Laboratory Dr. D. Gregory Arnold 2241 Avionics Circle Building 620 DSN: 785-1115 W-PAFB, OH 45433-7321				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/SNAT		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) N/A		
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT <p>Automatic target recognition (ATR) using radar commonly relies on modeling a target as a collection of point scattering centers. Features extracted from these scattering centers for input to a target classifier may be constructed that are invariant to translation and rotation in that they are independent of the position and aspect angle of the target in the radar scene.</p> <p>Here an iterative approach to efficiently build scattering center models is developed, and the shape space of these models is investigated. Experimental results are obtained for sets of three-dimensional scattering centers compressed to nineteen-dimensional feature sets, each consisting of the singular values of the matrix of scattering centers augmented with its second and third order monomial expansions. These feature sets are invariant to rotation and translation and permit the comparison of targets modeled by different numbers of scattering centers. A Mahalanobis distance metric is used that effectively identifies targets of military interest under the "real world" conditions that include noise and obscuration.</p> <p>In particular, eight targets of military interest are sampled in tenth-degree aspect angle increments to extract scattering centers, and 36 subclasses that encompass ten degrees are specified for each target. Each subclass is compressed to a nineteen-dimensional singular value feature set, and because the spatial distribution of the 100 nineteen-dimensional points in each subclass is approximately Gaussian, a mean and a covariance matrix represent each subclass. An unknown target is represented as point in the nineteen-dimensional feature space and matched to the closest subclass in Mahalanobis distance (distance in the direction of the principal eigenvector of the subclass covariance matrix). It is found that noisy targets are matched exactly in both target identity and aspect angle (to within ten degrees), whereas targets obscured by 20% are identified correctly in 80% of the test cases.</p> <p>This research has developed a practical means (1) to build scattering center models, (2) to compress each scattering center set into a small set of features for target classification, and (3) to implement classifiers that effectively function in the presence of noise and obscuration for targets of military interest.</p>						
15. SUBJECT TERMS target recognition, invariance, obscuration, Gaussian noise, eigenvalues						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Steven C. Gustafson, Associate Professor, USAF (ENG)	
U	U	U	UU	106	19b. TELEPHONE NUMBER (Include area code) (937) 255-3636, ext. 4598; e-mail: steven.gustafson@afit.edu	